arXiv:1402.6889v1 [cs.LO] 27 Feb 2014

Lazy Model Expansion: Interleaving Grounding with Search

Broes De Cat*, Marc Denecker, Peter Stuckey and Maurice Bruynooghe

February 28, 2014

Abstract

Finding satisfying assignments for the variables involved in a set of constraints can be cast as a (bounded) *model generation* problem: search for (bounded) models of a theory in some logic. The state-of-the-art approach for bounded model generation for rich knowledge representation languages, like Answer Set Programming (ASP), FO($\cdot$) and Zinc, is *ground-and-solve*: reduce the theory to a ground or propositional one and apply a search algorithm to the resulting theory.

An important bottleneck is the blowup of the size of the theory caused by the reduction phase. *Lazily grounding* the theory during search is a way to overcome this bottleneck. We present a theoretical framework and an implementation in the context of the FO($\cdot$) knowledge representation language. Instead of grounding all parts of a theory, *justifications* are derived for some parts of it. Given a partial assignment for the grounded part of the theory and valid justifications for the formulas of the non-grounded part, the justifications provide a recipe to construct a complete assignment that satisfies the non-grounded part. When a justification for a particular formula becomes invalid during search, a new one is derived; if that fails, the formula is split in a part to be grounded and a part that can be justified.

The theoretical framework captures existing approaches for tackling the grounding bottleneck such as lazy clause generation, grounding-on-the-fly and presents a generalization of the 2-watched literal scheme. We present an algorithm for lazy model expansion and integrate it in a model generator for FO(ID), a language extending first-order logic with inductive definitions. The algorithm is implemented as part of the state-of-the-art FO(ID) Knowledge-Base System IDP. Experimental results illustrate the power and generality of the approach.

*Broes De Cat is funded by the Agency for Innovation by Science and Technology in Flanders (IWT).

1 Introduction

For more and more search problems, *declarative modeling* is the method of choice to solve them: the problem is described in a declarative (logic) language in such a way that structures that satisfy the theory are solutions to the problem. Given such a theory and a (partial) input structure describing a particular problem instance that least specifies the domain, a *model generator* searches for a (bounded) structure that extends the input structure and satisfies the theory. Such a structure is a solution for the instance at hand. This task is known as *(bounded) model generation*. Model generation, studied for example in the field of Knowledge Representation (KR) (Baral, 2003), is closely related to the task of solving constraint satisfaction problems, studied in Constraint Programming (CP) (Apt, 2003), and that of generating answer sets of logic programs, studied in ASP (Marek & Truszczyński, 1999).

The aim of research into declarative modeling is to provide model generators for rich languages that allow users to naturally model domain-specific knowledge. To that end, state-of-the-art approaches often follow a two-phase methodology. In a first phase, the input theory, in the rich language at hand, is reduced into a fragment of the language that is supported by some search algorithm. In the second phase, the search algorithm is applied to the reduced theory to effectively search for models. For example, model generation for the language MiniZinc (Nethercote, Stuckey, Becket, Brand, Duck, & Tack, 2007) is performed by reducing to the ground language FlatZinc, for which search algorithms are available. Similarly, the language FO($\cdot$) (Denecker & Ternovska, 2008a) is reduced to its propositional fragment PC($\cdot$), for example as in (Wittocx, Mariën, & Denecker, 2010a), and ASP is reduced to propositional ASP, as in (Gebser, Schaub, & Thiele, 2007). As the reduced theory is often in a ground fragment of the language, we refer to the resulting reduced theory as the *grounding* and to the first phase as the *grounding* phase (where quantifiers are instantiated with elements in the domain). In other fields, grounding is also referred to as flattening, unrolling or splitting or propositionalization. The methodology itself is generally referred to as *ground-and-solve*.

Grounding becomes a bottleneck as users turn to applications with large domains and complex constraints. Indeed, it is easy to see that the grounding size of an FO formula is exponential in the nesting depth of quantifiers and in the arity of predicates and polynomial in the size of the universe of discourse. There is an increasing number of applications where the size of the grounded theory is so large that it does not fit in memory.

In this paper, we present a novel approach to remedy this bottleneck, called *lazy model expansion*, where the grounding is generated lazily (on-the-fly) during search, instead of up-front. The approach works by associating *justifications* to the non-ground parts of the theory. A valid justification for a non-ground formula is a recipe to expand a partial structure into a more precise (partial) structure that satisfies the formula. Given a partial structure and a valid justification for each of the non-ground formulas, a (total) structure can be constructed that satisfies all the non-ground formulas and extends the given structure. Consequently, model generation can be limited to the grounded part of the theory; if a model is found for that part, it can be extended to a model of the whole theory. However, a new assignment during model generation can conflict with one of the justifications. In that case, an alternative justification needs to be sought. If none is found, the associated formula can be split in two parts, one part that is grounded and one part for which a valid justification is still available.

Example 1.1. Consider the `Sokoban` toy problem, a planning problem where a robot has to push blocks around on a 2-D grid to arrange them in a given goal configuration. A constraint on the move action is that the target position of the moved block is currently empty, which can be expressed as

$$\forall(t, b, p) \in time \times block \times pos : move(b, p, t) \Rightarrow empty(p, t). \qquad (1)$$

As it is not known in advance how many time steps are needed, one ideally wants to assume a very large or even infinite number of steps. Using ground-and-solve, this blows-up the size of the grounding.

Given a partial plan P (a structure) that only moves block b_1 to position p_1 at time t_1 and in which $empty(p_1, t_1)$ is true, the partial plan does not violate this constraint. In that case, there is a simple way to extend the plan so that it certainly satisfies the constraint: do not move any more blocks. Indeed, the recipe "$move(b, p, t)$ is false except for $move(b_1, p_1, t_1)$" can be used to *construct* a complete plan in which sentence 1 is satisfied.

We do not immediately apply this recipe, but use it as a *default*, to be applied after a (partial) plan has been found that satisfies the goal configuration. However, if during search it is decided to move a block b_2 at time t_2 to position p_2, this conflicts with the recipe, so we are no longer guaranteed it will result in a complete, valid plan. In that case, one can split up sentence (1) into sentences (2), for the instantiation $\langle time/t_2, block/b_2, pos/p_2\rangle$, and (3), for the remainder. For sentence (3), if $move(b_2, p_2, t_2)$ is added as (second) exception, the recipe still holds; hence only sentence (2) needs to

be grounded and taken into account during search.

$$move(b_2, p_2, t_2) \Rightarrow empty(p_2, t_2) \quad (2)$$

$$\forall (t, b, p) \in time \times block \times pos \backslash (t_2, b_2, p_2) : move(b, p, t) \Rightarrow empty(p, t) \quad (3)$$

Lazy grounding is a form of constraint propagation. In general, a constraint propagator for a constraint c is responsible for detecting that an assignment to variables in c violates the constraint and to derive consequences of c under (partial) assignments. Either by observing that a particular variable is forced to take a specific value or by adding extra constraints to the constraint store. Lazy grounding acts as a propagator for non-ground constraints. As soon as a constraint's justification is violated, it adds some new ground constraints to the store and creates/reuses justifications for the remaining non-ground parts of the constraint. Interestingly, it can be shown (see Section 5) that justifications can be chosen in such a way that propagation that would occur in the fully grounded theory is not delayed.

The main contributions of this paper are:

- A theoretical framework for *lazy model expansion*. By aiming at minimally instantiating quantified variables, it paves the way for a solution to the long-standing problem of handling quantifiers in search problems, encountered, e.g., in the fields of ASP (Lefèvre & Nicolas, 2009) and SAT Modulo Theories (Ge & de Moura, 2009). The framework also generalizes existing approaches that are related to the grounding bottleneck such as incremental domain extension (Claessen & Sörensson, 2003) and lazy clause generation (Stuckey, 2010).

- A complete algorithm for lazy model expansion for the logic FO(ID), the extension of first-order logic (FO) with inductive definitions (a language closely related to ASP as shown in (Denecker, Lierler, Truszczynski, & Vennekens, 2012)). This includes efficient algorithms to derive consistent sets of justifications and to maintain them throughout changes in a partial structure (e.g., during search).

- An implementation extending the IDP knowledge-base system (De Cat, Bogaerts, Bruynooghe, & Denecker, 2014) and experiments that illustrate the power and generality of lazy grounding.

The paper is organized as follows. In Section 2, the necessary background and notations are introduced. Formal definitions of the concepts used for lazy grounding are presented in Section 3, followed by a presentation of the relevant algorithms and optimizations in Sections 4 and 5. Experimental

evaluation is provided in Section 6, followed by a discussion on related and future work and a conclusion. A preliminary version of this work appeared as (De Cat, Denecker, & Stuckey, 2012).

2 Preliminaries

2.1 FO(ID)

This paper makes use of FO(ID), a logic that extends FO with inductive definitions (Denecker & Ternovska, 2008a). This section introduces some notations and elaborates on what is needed beyond the basics of FO.

Without loss of generality, we limit FO and FO(ID) to its function-free fragment. Function symbols can always be replaced by graph predicates (Enderton, 2001).

A vocabulary Σ of the function-free fragment consists of a set of predicate symbols. Propositional symbols are 0-ary predicate symbols; these include the symbols $\top$ and $\bot$, denoting true, respectively false. Predicate symbols are usually denoted by P, Q and R; atoms by a and literals (atoms or their negation) by l. Variables are denoted by x and y, domain elements by d, an ordered set of elements $e_1, \ldots, e_n$ by $\bar{e}$.

In the context of this paper, theories are considered in the context of a given domain D and we use formulas that contain elements of D. A *domain atom* is an atom of the form $P(\overline{d})$ with $\overline{d}$ an n-tuple of domain elements and P a predicate symbol of arity n. Likewise, we consider *domain literals*.

A structure $\mathcal{I}$ interpreting Σ consists of a set of domain elements D (the *domain*) and an n-ary relation $P^{\mathcal{I}}$ for all predicate symbols $P \in \Sigma$ of arity n. The propositional symbols $\top$ and $\bot$ are respectively interpreted as $\mathbf{t}$ and $\mathbf{f}$. We view an n-ary relation as a function $D^n \to \{\mathbf{t}, \mathbf{f}\}$. Model generation algorithms maintain *partial* structures and may (temporarily) find themselves in an *inconsistent* state, for example when a conflict arises. To represent such states, three-valued (or partial) and four-valued structures are introduced, which interpret n-ary predicate symbols as functions $D^n \to \{\mathbf{t}, \mathbf{f}, \mathbf{u}, \mathbf{i}\}$. A structure is determined by its domain and the truth value of its domain atoms. If D and Σ are fixed, we can use the following alternative way to represent a four-valued structure with domain D, namely as a set of domain literals. Indeed, there is a one-to-one correspondence between such sets S and four-valued structures: for a domain atom a, $a^S = \mathbf{i}$ (inconsistent) if both a and $\neg a$ are in S, $a^S = \mathbf{t}$ if only a is in S, $a^S = \mathbf{f}$ if only $\neg a$ is in S and $a^S = \mathbf{u}$ (unknown) otherwise. From now on, we may treat four-valued structures as sets of domain literals and vice versa.

A structure $\mathcal{I}$ interpreting Σ naturally corresponds to structure $\mathcal{I}'$ interpreting a vocabulary $\Sigma' \supset \Sigma$, namely by defining $a^{\mathcal{I}'} = a^{\mathcal{I}}$ if a is a domain atom using a predicate of Σ and $a^{\mathcal{I}'} = \mathbf{u}$ otherwise. Note that both $\mathcal{I}$ and $\mathcal{I}'$ correspond to the same set of domain literals.

A *two-valued* structure has range $\{\mathbf{t}, \mathbf{f}\}$ for its relations and is a special case of a *partial* or *three-valued* structure (range $\{\mathbf{t}, \mathbf{f}, \mathbf{u}\}$) which itself is a kind of four-valued structure (range $\{\mathbf{t}, \mathbf{f}, \mathbf{u}, \mathbf{i}\}$). A structure is inconsistent if $\mathbf{i}$ is in the image of the interpretation of at least one predicate symbol. When unqualified, the term *structure* stands for a four-valued structure.

For any set of domain atoms S, we use $\mathcal{I}|_S$ to denote the restriction of $\mathcal{I}$ to S: $a^{\mathcal{I}|_S} = a^{\mathcal{I}}$ if $a \in S$ and $a^{\mathcal{I}|_S} = \mathbf{u}$ otherwise. If S is a set of predicate symbols, then $\mathcal{I}|_S$ is the restriction of $\mathcal{I}$ to the set of all domain atoms of these symbols. We call $\mathcal{I}$ a two-valued structure of S if $\mathcal{I}$ is two-valued on domain atoms of S and unknown otherwise.

The (partial) *truth* order $>_t$ on truth values is defined by $\mathbf{t} >_t \mathbf{u} >_t \mathbf{f}$ and $\mathbf{t} >_t \mathbf{i} >_t \mathbf{f}$. The (partial) *precision* order $>_p$ on truth values is defined by $\mathbf{i} >_p \mathbf{t} >_p \mathbf{u}$ and $\mathbf{i} >_p \mathbf{f} >_p \mathbf{u}$. For two structures $\mathcal{I}$ and $\mathcal{I}$' with domain D interpreting the same vocabulary, we say that $\mathcal{I}$' is an *expansion* of $\mathcal{I}$ if $P^{\mathcal{I}'}(\overline{d}) \geq_p P^{\mathcal{I}}(\overline{d})$ for all predicate symbols P and tuples of domain elements $\overline{d}$. Viewing structures as sets of domain literals, this corresponds to $\mathcal{I} \subseteq \mathcal{I}'$.

We assume familiarity with the syntax of (function-free) FO. The main difference with standard FO is that quantification is over explicitly specified subsets of the domain D and is denoted $\exists x \in D' : \varphi$ and $\forall x \in D' : \varphi$, with $D' \subseteq D$. We sometimes abbreviate $\exists x_1 \in D_1 : \ldots \exists x_n \in D_n$ as $\exists \overline{x} \in \overline{D}$, and similarly for $\forall$. Given a formula φ, $\varphi[\overline{x}]$ indicates that $\overline{x}$ are the free variables of φ. Substitution of a variable x in formula φ by a term t is denoted by $\varphi[x/t]$. A *ground formula* (in domain D) is a formula without variables (hence without quantifiers). Similar properties and notations are used for *rules* (introduced below).

We denote by $voc(T)$ the set of all predicate symbols that occur in theory T. For a structure $\mathcal{I}$, $voc(\mathcal{I})$ is the set of symbols interpreted by $\mathcal{I}$. Unless specified otherwise, theories and structures range over the vocabulary Σ.

The language FO(ID) extends FO with (inductive) *definitions*. A definition Δ is a set of rules of the form $\forall \overline{x} \in \overline{D} : P(x_1, \ldots, x_n) \leftarrow \varphi$, with P a predicate symbol and φ an FO formula. The literal $P(\overline{x})$ is referred to as the *head* of the rule and φ as the *body*. Given a rule r, the $head(r)$ and $body(r)$ denote respectively the head and the body of r. Given a definition Δ, a domain atom $P(\overline{d})$ is *defined* by Δ if there exists a rule $\forall \overline{x} \in \overline{D} : P(\overline{x}) \leftarrow \varphi$ in Δ such that $\overline{d} \in \bar{D}$. Otherwise $P(\overline{d})$ is *open* in Δ. A domain literal $[\neg]P(\overline{d})$

is defined by Δ if $P(\overline{d})$ is defined by Δ. The sets of defined and open domain atoms of Δ are denoted as $defined(\Delta)$ and $open(\Delta)$, respectively. A theory in FO(ID) consists of a set of sentences and a set of definitions.

Without loss of generality, we assume that in any definition a domain atom is defined by at most one rule.

2.1.1 Model semantics

In this paper, we use the standard two-, three- and four-valued truth assignment function, defined by structural induction for pairs of FO domain formulas φ and structures $\mathcal{I}$ that interpret φ:

- $P(\overline{d})^{\mathcal{I}} = P^{\mathcal{I}}(\overline{d}^{\mathcal{I}})$,
- $(\psi \wedge \phi)^{\mathcal{I}} = min_{<_t}(\psi^{\mathcal{I}}, \phi^{\mathcal{I}})$,
- $(\psi \vee \phi)^{\mathcal{I}} = max_{<_t}(\psi^{\mathcal{I}}, \phi^{\mathcal{I}})$,
- $(\neg\psi)^{\mathcal{I}} = (\psi^{\mathcal{I}})^{-1}$,
- $(\exists x \in D : \psi)^{\mathcal{I}} = max_{<_t}(\{\psi[x/d]^{\mathcal{I}} \mid d \in D\})$,
- $(\forall x \in D : \psi)^{\mathcal{I}} = min_{<_t}(\{\psi[x/d]^{\mathcal{I}} \mid d \in D\})$.

Monotonicity holds: if $\mathcal{I} \leq_p \mathcal{I}'$, then $\varphi^{\mathcal{I}} \leq_p \varphi^{\mathcal{I}'}$. Hence, if a formula is true in a partial structure, it is true in all two-valued expansions of it.

A structure $\mathcal{I}$ is a *model* of / *satisfies* a sentence φ (notation $\mathcal{I} \models \varphi$) if $\mathcal{I}$ is two-valued and $\varphi^{\mathcal{I}} = \mathbf{t}$. This satisfaction relation can be extended to handle definitions by means of the well-founded semantics (Van Gelder, 1993). This semantics formalizes the informal semantics of rule sets as inductive definitions (Denecker, 1998; Denecker, Bruynooghe, & Marek, 2001). An interpretation $\mathcal{I}$ satisfies a definition Δ ($\mathcal{I} \models \Delta$) if $\mathcal{I}$ is two-valued and interprets $defined(\Delta)$ as $wf_{\Delta}(\mathcal{I}|_{open(\Delta)})$, the well-founded model of Δ in the structure $\mathcal{I}|_{open(\Delta)}$(Van Gelder, 1993; Denecker & Ternovska, 2008b). A structure $\mathcal{I}$ satisfies a theory $\mathcal{T}$ if $\mathcal{I}$ is two-valued and $\mathcal{I}$ is a model of all sentences and definitions in $\mathcal{T}$. Below, we present a formalization of the well-founded semantics using the notion of *justification*.

A definition Δ is *total* if for each structure $\mathcal{I}$ that is two-valued on $open(\Delta)$, a two-valued expansion of $\mathcal{I}$ exists which is a model of Δ (Denecker & Ternovska, 2008b). Equivalently, this is the case if the well-founded model of Δ in $\mathcal{I}$ is two-valued. Broad classes of definitions have been proven to be total (Denecker & Ternovska, 2008b). As explained in the same paper, it is

a methodological guideline in FO(ID) that the user write total definitions. The techniques presented below require that definitions are total.

Two theories T and T', possibly over different vocabularies, are Σ-*equivalent* if each model of T restricted to Σ can be expanded to a model of T' and vice versa. Two theories T and T' are *strongly* Σ-*equivalent* if the above expansions are also unique. By extension, (strong) Σ-equivalence *in a structure* $\mathcal{I}$ is defined similarly: if each model of T expanding $\mathcal{I}$ can be expanded to a model of T' expanding $\mathcal{I}$ and vice versa; to obtain strong equivalence, these expansions have to be unique. From a theory T, we often derive a *strongly* $voc(T)$-equivalent theory T' in a given structure $\mathcal{I}$. Such transformations preserve satisfiability and number of models *and* each model of T' can be directly mapped to a model of T by projection on $voc(T)$.

Canonical theories. To simplify the presentation, the lazy grounding techniques are presented here for theories of the form $\{P_{\mathcal{T}}, \Delta\}$, with $P_{\mathcal{T}}$ a propositional symbol, and Δ a single definition with function-free rules. This is without loss of generality. First, as mentioned above, standard techniques allow to make a theory function-free. Second, multiple definitions can be combined into one as described in (Denecker & Ternovska, 2008b; Mariën, Gilis, & Denecker, 2004), for example if a stratification exists or by introduction of additional predicate symbols. Furthermore, we assume negation only occurs in front of atoms and compound formulas are built by conjunction, disjunction and (existential and universal) quantification. We also assume that D is the domain of structures.

The methods below can be extended to standard theories with functions, FO axioms and multiple definitions, and such extended methods have been implemented in our system. However, this introduces a number rather irrelevant technicalities which we want to avoid here.

2.1.2 Justifications

Recall that structures correspond one-to-one to sets of domain literals and that, by assumption, for each defined domain atom $P(\overline{d})$ there is a unique rule $\forall \overline{x} \in \overline{D} : P(\overline{x}) \leftarrow \varphi \in \Delta$ such that $\overline{d} \in \overline{D}$.

Definition 2.1 (Direct justification)**.** A *direct justification* for a defined domain literal $P(\overline{d})$ (respectively $\neg P(\overline{d})$) is a consistent set S of domain literals such that for the rule $\forall \overline{x} \in \overline{D} : P(\overline{x}) \leftarrow \varphi$ of Δ such that $\overline{d} \in \overline{D}$, it holds that $\varphi[\overline{x}/\overline{d}]^S = \mathbf{t}$ (respectively $\varphi[\overline{x}/\overline{d}]^S = \mathbf{f}$).

We observe that a direct justification is never empty since all formulas are unknown in the structure that assigns **u** to all domain atoms (including $\top$ and $\bot$).

Definition 2.2 (Justification)**.** A *justification* over Δ is a graph J on the set of domain literals such that if $(l, l') \in J$, then $\{l'' \mid (l, l'') \in J\}$ is a direct justification of l.

With a justification J, we can associate the total function from domain literals l to sets $J(l) = \{l'' \mid (l, l'') \in J\}$ of domain literals. If $J(l) \neq \emptyset$, then l is a defined literal and $J(l)$ is a direct justification of l. In the sequel we say that J is defined in l if $J(l) \neq \emptyset$. A justification is denoted as a set of pairs $l \to S$, with S the direct justification of l.

Example 2.3. Consider a domain $D = \{d_1, \ldots, d_n\}$ and the definition Δ

$$\left\{ \begin{array}{ll} \forall x \in D : P(x) & \leftarrow Q(x) \vee R(x) \\ \forall x \in D : Q(x) & \leftarrow P(x) \end{array} \right\}$$

The unique (minimal) direct justification for $Q(d_i)$ is $\{P(d_i)\}$ and for $\neg Q(d_i)$ is $\{\neg P(d_i)\}$. Direct justifications for $P(d_i)$ are both $\{Q(d_i)\}$ and $\{R(d_i)\}$ while the only (minimal) direct justification for $\neg P(d_i)$ is $\{\neg Q(d_i), \neg R(d_i)\}$. Atoms $R(d_i)$ are open and have no direct justification.

Definition 2.4 (Justification subgraph)**.** Let J be a justification over Δ. The justification *for a literal* l is the subgraph J_l of nodes and edges of J reachable from l. The justification *for a set of literals* L is the subgraph J_L of nodes and edges of J reachable from any $l \in L$.

A justification J is *total* for l if J is defined in each defined literal reachable from l; it is total for a set of literals L if it is total for each literal in L; it is total if it is total for all literals for which J is defined. A justification J is *consistent with* a structure $\mathcal{I}$ when none of the literals in J is false in $\mathcal{I}$.

A path in a justification J is a sequence $l_0 \to l_1 \to \ldots$ such that, if $l_i \to l_{i+1}$, then there is an edge from l_i to l_{i+1} in J. An infinite path may be cyclic or not.

Definition 2.5. A *cycle* in a justification J is the set of domain literals on an infinite path of J. A cycle is *positive* if the path has a tail consisting of positive literals; it is *negative* if the path has a tail consisting of negative literals; it is *mixed* otherwise.

Cycles are finite iff the path is cyclic. If the domain D is finite, cycles are always finite. Intuitively, a justification provides an argument to the truth of its nodes. If this argument contains a positive cycle, we consider it unfounded; negative cycles, though, are accepted.

Definition 2.6 (Justifies)**.** A justification J over Δ *justifies* a set of literals L defined in Δ (a set of literals *has a justification*) if (i) J_L is total for L; (ii) cycles in J_L, if any, are negative; (iii) the set of literals in J_L is consistent.

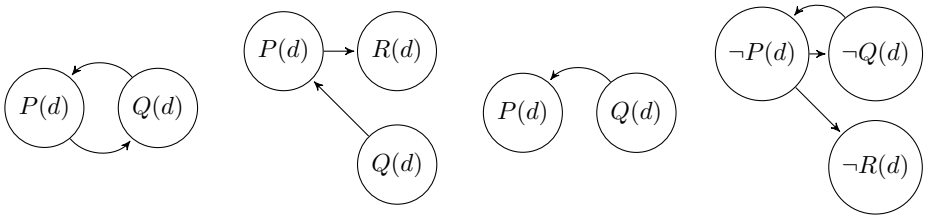

Figure 1: Justifications for definition Δ in Example 2.3, with $d \in D$.

Example 2.7. In Figure 1, we show a few possible justifications (ordered (i)-(iv) from left to right) over definition Δ in Example 2.3 that contain the defined domain atoms $P(d)$ and $Q(d)$ ($d \in D$). Justification (ii) justifies $Q(d)$ and (iv) justifies $\neg Q(d)$; (iii), however, is not total for $Q(d)$ and (i) has a positive cycle.

The relationship between justifications and the well-founded semantics has been investigated in different publications (Denecker & De Schreye, 1993, 1992; Mariën, 2009). Below we recall the results on which this paper relies. The first result states that if J justifies all literals in L, then any model $\mathcal{I}$ of Δ in which the leaves of J_L are true, satisfies all literals in L and in J_L.

Proposition 2.8. *If J is a justification over Δ that justifies a set of domain literals L then all literals in J_L are true in every model of Δ in which the (open) leaves of J_L are true.*

For an interpretation $\mathcal{I}_{open}$ that is two-valued for $open(\Delta)$ an in which the leaves are true, the structure $\mathcal{I} = wf_\Delta(\mathcal{I}_{open})$ can be computed in time polynomial in the size of the domain, as shown in (Chen & Warren, 1996). When Δ is total, $wf_\Delta(\mathcal{I}_{open})$ is the well-founded model of Δ; otherwise, $wf_\Delta(\mathcal{I}_{open})$ can be a partial structure.

Example 2.9 (Continued from Example 2.7)**.** Consider justification (ii), with $R(d)$ the only literal open in the subgraph for $L = \{Q(d)\}$. Take $\mathcal{I}_{open}$

to be the structure that makes $R(d)$ true and all other domain atoms over R false. Then $\mathcal{I}_{open}$ is an $open(\Delta)$-structure $\geq_p \{R(d)\}$. The well-founded evaluation assigns $P(d)$ true (as $R(d)$ is true) and afterwards assigns $Q(d)$ true (as $P(d)$ is now true). Hence, $wf_\Delta(\mathcal{I}_{open})$ is a two-valued model of Δ (Δ is total) and expands L and all literals reachable from L.

Proposition 2.10. *If $\mathcal{I}$ is a model of Δ, then there exists a justification J over Δ that justifies each defined domain literal true in $\mathcal{I}$ and contains only domain literals true in $\mathcal{I}$.*

Corollary 2.11. *In case Δ is total, if a justification J over Δ justifies a set of domain literals L, then every two-valued $open(\Delta)$-structure consistent with J_L can be extended in a unique way to a model of Δ that satisfies all literals of L.*

Hence, for a canonical theory $\{P_\mathcal{T}, \Delta\}$ (recall, Δ is total), the theory is satisfiable iff a justification J exists that justifies $P_\mathcal{T}$.

2.2 Generating models

Model generation is the inference task that takes as input a theory $\mathcal{T}$ and returns as output a model of $\mathcal{T}$. Model Expansion (MX) was defined in (Mitchell, Ternovska, Hach, & Mohebali, 2006) as the inference task that takes as input a theory $\mathcal{T}$ over vocabulary Σ and a two-valued structure $\mathcal{I}$ over a subvocabulary of Σ, and returns an expansion $\mathcal{M}$ of $\mathcal{I}$ satisfies $\mathcal{T}$. Here, it will be the more general inference problem as defined in (Wittocx, Mariën, & Denecker, 2008) that takes as input a (potentially partial) structure $\mathcal{I}$ over Σ, and returns an expansion $\mathcal{M}$ of $\mathcal{I}$ that satisfies $\mathcal{T}$.

As already mentioned, the state-of-the-art approach to model expansion solvers in FO(ID) is, similar as in ASP, grounding $\mathcal{T}$ in the context of $\mathcal{I}$ and afterwards applying search to the resulting ground theory. An example of the latter sort of algorithm is the SAT(ID) search algorithm (Mariën, Wittocx, Denecker, & Bruynooghe, 2008).

Below, we present the grounding algorithm that will be the basis of the lazy MX algorithm. We assume familiarity with the basic Conflict-Driven Clause-Learning (CDCL) algorithm of SAT algorithms (Mitchell, 2005).

2.2.1 Grounding

For an overview of intelligent grounding techniques in FO(ID), we refer the reader to (Wittocx, Denecker, & Bruynooghe, 2013a) and (Wittocx et al., 2010a). Below we present the basic principle.

A grounder takes as input a theory $\mathcal{T}$ over vocabulary Σ, a partial structure $\mathcal{I}$ with domain D, interpreting at least $\top$ and $\bot$, and returns a ground theory $\mathcal{T}'$ that is strongly Σ-equivalent with $\mathcal{T}$ in $\mathcal{I}$. Theory $\mathcal{T}'$ is then called a *grounding* of $\mathcal{T}$ given $\mathcal{I}$. Recall that we assume that $\mathcal{T}$ is a canonical theory of the form $\{P_\mathcal{T}, \Delta\}$.

One way to compute the grounding is using a top-down process on the theory, iteratively applying grounding steps to direct subformulas of the rule or formula at hand. The grounding algorithm may replace subformulas by new predicate symbols as follows. Let $\varphi[\overline{x}]$ be a formula in $\mathcal{T}$ and let $\overline{D}$ be the domains of $\overline{x}$. *Tseitin transformation* replaces φ by the atom $T_\varphi(\overline{x})$, with T_φ a new $|\overline{x}|$-ary predicate symbol called a *Tseitin* symbol, and extends Δ with the rule $\forall \overline{x} \in \overline{D} : T_\varphi(\overline{x}) \leftarrow \varphi$. The new theory is strongly Σ-equivalent to the original one (Vennekens, Mariën, Wittocx, & Denecker, 2007).

The procedure one_step_ground, outlined in Figure 1, performs one step in the grounding process. Called with a formula or rule φ in canonical form, the algorithm replaces all direct subformulas with Tseitin symbols and returns a pair consisting of a ground part G (rules or formulas) and a possibly non-ground part R (rules). If φ is a formula, then G consists of ground formulas. Replacing φ by the returned ground formulas and extending Δ with the returned rules produces a theory that is strongly $voc(\mathcal{T})$-equivalent to the original. If φ is a rule from Δ, it is replaced by both sets of returned rules and again, the new theory is strongly $voc(\mathcal{T})$-equivalent to the original.

Grounding a theory then boils down to applying one_step_ground on the sentence $P_\mathcal{T}$ (which copies $P_\mathcal{T}$ to the ground part) and on each rule of the theory and repeatedly applying one_step_ground on the returned rules R (all returned sentences and rules in G are ground). We use ground to refer to the algorithm for this overall process.

Various improvements exist, such as returning $\top/\bot$ for atoms interpreted in $\mathcal{I}$ and returning $\bot$ from conjunctions whenever a false conjunct is encountered (analogously for disjunctions and quantifications).

Also, algorithm one_step_ground introduces a large number of Tseitin symbols. State-of-the-art grounding algorithms use a number of optimizations to reduce the number of such symbols. As these optimizations are not directly applicable to the techniques presented in this paper, we start from the naive one_step_ground algorithm. In Section 5, we present an optimized version of one_step_ground that introduces less Tseitin symbols and hence results in smaller groundings.

Algorithm 1: The one_step_ground algorithm.

```
1  Function one_step_ground (formula or rule φ)
2      switch φ do
3          case [¬]P(d̄) return ⟨{φ}, ∅⟩;
4          case P(d̄) ← ψ
5              ⟨G, Δ⟩ := one_step_ground(ψ);
6              return ⟨{P(d̄) ← ⋀_{g∈G} g}, Δ⟩;
7          case ψ_1 ∨ ... ∨ ψ_n
8              return ⟨{⋁_{i∈[1,n]} T_{ψ_i}}, {T_{ψ_i} ← ψ_i | i ∈ [1,n]}⟩;
9          case ψ_1 ∧ ... ∧ ψ_n
10             return ⟨{T_{ψ_i} | i ∈ [1,n]}, {T_{ψ_i} ← ψ_i | i ∈ [1,n]}⟩;
11         case ∀x̄ ∈ D̄ : P(x̄) ← ψ
12             return ⟨∅, {P(x̄)[x̄/d̄] ← ψ[x̄/d̄] | d̄ ∈ D̄}⟩;
13         case ∃x̄ ∈ D̄ : ψ[x̄]
14             return ⟨{⋁_{d∈D} T_{ψ[x̄/d̄]}}, {T_{ψ[x̄/d̄]} ← ψ[x̄/d̄] | d̄ ∈ D̄}⟩;
15         case ∀x̄ ∈ D̄ : ψ[x̄]
16             return ⟨{T_{ψ[x̄/d̄]} | d̄ ∈ D̄}, {T_{ψ[x̄/d̄]} ← ψ[x̄/d̄] | d̄ ∈ D̄}⟩;
```

3 Lazy grounding and lazy model expansion

We use the term *lazy grounding* to refer to the process of partially grounding a theory and the term *lazy model expansion* (lazy MX) for the process that interleaves lazy grounding with model expansion over the grounded part. In Section 3.1, we formalize a framework for lazy model expansion of FO(ID) theories; in Section 3.2, we formalize the instance of this framework that is the basis of our current implementation; in Section 3.3, we illustrate its operation.

3.1 Lazy model expansion for FO(ID) theories

The input of the lazy MX algorithm consists of a canonical theory $\mathcal{T} = \{P_\mathcal{T}, \Delta\}$ and an input structure $\mathcal{I}_{in}$ with domain D interpreting at least $\top$ and $\bot$. We assume that Δ is total (in $\mathcal{I}_{in}$). The task is then to find models of Δ in which $P_\mathcal{T}$ is true. This is achieved by interleaving lazy grounding with search on the already grounded part. We first focus on the lazy grounding.

The initial input to lazy grounding consists of $P_\mathcal{T}$ (handled separately), Δ and $\mathcal{I}_{in}$. For subsequent steps, the input consists of a set of rules still to

be grounded, an already grounded theory and a three-valued structure that is an expansion of the initial input structure.

Each subsequent grounding step can replace non-ground rules by ground rules and might introduce new rules. Hence, the state of the grounding includes a set Δ_{g} of ground rules and a set Δ_{d} of (possibly) non-ground rules. The definitions have the property that $\Delta_{\mathrm{g}} \cup \Delta_{\mathrm{d}}$ (in what follows abbreviated as Δ_{gd}) is voc(Δ)-equivalent with the original definition Δ and hence, Δ_{gd} is total. The grounding procedure will guarantee that, at all times, Δ_{g} and Δ_{d} are total.

Given a partial structure $\mathcal{I}_{in}$ and the rule sets Δ_{g} and Δ_{d}, the key idea behind lazy model expansion is (**i**) to use a search algorithm to search for a model $\mathcal{I}$ of Δ_{g} that is an expansion of $\mathcal{I}_{in}$ in which $P_{\mathcal{T}}$ is true; (**ii**) to maintain a justification J such that the literals true in $\mathcal{I}$ and defined in Δ_{d} are justified over Δ_{gd} and that J is consistent with $\mathcal{I}$; (**iii**) to interleave steps (**i**) and (**ii**) and to move parts of Δ_{d} to Δ_{g} when some literal defined in Δ_{d} that needs to be justified cannot be justified.

Thus, to control lazy model expansion, it suffices to maintain a state $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ consisting of the grounded rules Δ_{g}, the rules Δ_{d} yet to be grounded, a justification J, and a three-valued structure $\mathcal{I}$. Initially, $\mathcal{I}$ is $\mathcal{I}_{in}$, Δ_{d} is Δ, $\Delta_{\mathrm{g}} = \emptyset$, and J is the empty graph.

Lazy model expansion searches over the space of *acceptable* states.

Definition 3.1 (Acceptable state)**.** A tuple $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ of a theory with an atomic sentence $P_{\mathcal{T}}$, a total definition Δ, and an input structure $\mathcal{I}_{in}$ is an *acceptable* state if (**i**) Δ_{gd}, Δ_{g} and Δ_{d} are total definitions and Δ_{gd} is strongly voc(Δ)-equivalent with Δ, (**ii**) no domain atom is defined in both Δ_{g} and Δ_{d}, (**iii**) J is a justification over Δ_{gd}, (**iv**) $\mathcal{I}$ is an expansion of $\mathcal{I}_{in}$, (**v**) the set L of literals true in $\mathcal{I}$ and defined in Δ_{d} are justified by J, and (**vi**) J_L is consistent with $\mathcal{I}$.

The lazy model expansion algorithm tries to compute an acceptable state in which $P_{\mathcal{T}}$ is justified. By Corollary 2.11, this would entail that a model of $\mathcal{T}$ exists; it can be computed efficiently through well-founded model computation. In intermediate states, the justification may be non-total for $P_{\mathcal{T}}$, contain positive or mixed cycles, or be inconsistent.

Note (**iii**) that the justification must be over Δ_{gd}. Indeed, assume some literal l is justified over Δ_{d}. Its justification graph can have a leaf that is defined in Δ_{g} and that depends positively or negatively on l. Then every attempt to extend this justification graph to a total justification graph that justifies l over Δ_{gd} might fail because of a forbidden cycle.

Proposition 3.2. *Let $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ be an acceptable state. Δ_{gd} has a well-founded model that expands the literals that are true in $\mathcal{I}$ and defined in the delayed definition Δ_d.*

Proof. Let L be the set of literals true in $\mathcal{I}$ and defined in Δ_{d}. As the state is acceptable, J justifies the literals of L. Hence, by Corollary 2.11, there exists a well-founded model that expands L. □

Example 3.3. Consider the theory $\{P_{\mathcal{T}}, \Delta\}$, with Δ the definition

$$\left\{ \begin{array}{l} P_{\mathcal{T}} \leftarrow T_1 \vee T_2 \vee T_3 \\ T_1 \leftarrow \forall x \in D : Q(x) \\ T_2 \leftarrow \forall x \in D : R(x) \\ T_3 \leftarrow \exists x \in D : \neg Q(x) \end{array} \right\}$$

Let $\mathcal{I}$ be the structure $\{P_{\mathcal{T}}, T_1\}$ (hence, T_2 and T_3 are unknown), and Δ_{g} and Δ_{d} the definitions consisting of the first rule, respectively the remaining rules. Furthermore, let J be $\{T_1 \rightarrow \{Q(d) \mid d \in D\}\}$. The tuple $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ is then an acceptable state. Indeed, T_1 is the only literal in $\mathcal{I}$ that is defined in Δ_{d} and it is justified by J. The well-founded evaluation, after assigning $\top$ to the open literals of J (i.e., to $\{Q(d) \mid d \in D\}$), derives that T_1 is true. Moreover, because $\mathcal{I}$ is a model of Δ_{g}, $P_{\mathcal{T}}$ is also true in such a well-founded model. Note that R can be interpreted randomly, as no R-atoms occur in $\mathcal{I}$ or J.

The following theorem states conditions on when the obtained expansion is also a model of $\mathcal{T}$.

Theorem 3.4. *Let $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ be an acceptable state of a theory $\mathcal{T} = (P_{\mathcal{T}}, \Delta)$ with input structure $\mathcal{I}_{in}$ such that $P_{\mathcal{T}}$ is true in $\mathcal{I}$ and $\mathcal{I}|_{voc(\Delta_g)}$ is a model of Δ_g. Then there exists a model $\mathcal{M}$ of $\mathcal{T}$ that expands $\mathcal{I}|_{voc(\Delta_g)}$.*

Proof. $\mathcal{I}|_{voc(\Delta_{\mathrm{g}})}$ is a model of Δ_{g}. It follows from Proposition 2.10 that there exists a justification J_g over Δ_{g} that justifies every true defined literal of Δ_{g} and that consists of only domain literals true in $\mathcal{I}|_{voc(\Delta_{\mathrm{g}})}$. We now have two justifications: J and J_g. We combine them in one J_c as follows: for each defined literal l of Δ_{gd}, if J is defined in l, we set $J_c(l) = J(l)$; otherwise, we set $J_c(l) = J_g(l)$.

We verify that J_c justifies $P_{\mathcal{T}}$. First, it is total in $P_{\mathcal{T}}$. Indeed, any path from $P_{\mathcal{T}}$ either consists of literals defined in Δ_{g}, and then it is a branch of the total J_g over Δ_{g}, or it passes to a literal l' defined in Δ_{d}, which is

justified by J according to condition (**v**) and hence $(J_c)_{l'} = J_{l'}$ is total. As such, from $P_{\mathcal{T}}$ we cannot reach a defined literal of Δ_{gd} in which J_c is undefined. Second, J_c does not contain mixed or positive cycles starting from $P_{\mathcal{T}}$. This is because any path from $P_{\mathcal{T}}$ is either a path in J_g or it has a tail in J, and neither of these contain mixed or positive cycles. Finally, the set of literals reachable from $P_{\mathcal{T}}$ in J_c is consistent. Also this we can see if we look at paths in J_c from $P_{\mathcal{T}}$: at first we follow J_g which consists of true literals in $\mathcal{I}$, then we may get into a path of J which contains literals that are consistent with $\mathcal{I}$. In any case, it is impossible to reach both a literal and its negation.

It follows from Proposition 2.8 that there exists a model of Δ_{gd} that expands $\mathcal{I}|_{voc(\Delta_{\text{g}})}$ and in which $P_{\mathcal{T}}$ is true. Since Δ_{gd} is strongly equivalent with Δ, the proposition follows. □

Recall that effectively computing such a model $\mathcal{M}$ can be achieved by well-founded evaluation of Δ_{gd}, with polynomial data complexity, starting from any two-valued $open(\Delta_{\text{gd}}) - structure$ expanding $\mathcal{I}|_{voc(\Delta_{\text{g}})}$.

In the above theorem, it is required that $\mathcal{I}$ is a model of Δ_{g}. Actually, we do not need to compute a two-valued model of Δ_{g}. It suffices to search for a partial structure and a justification that justifies $P_{\mathcal{T}}$. So, we can relax this requirement at the expense of also maintaining justifications for literals true in $\mathcal{I}$ and defined in Δ_{g}.

Corollary 3.5. *Let $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ be an acceptable state of a theory $\mathcal{T} = \{P_{\mathcal{T}}, \Delta\}$ with input structure $\mathcal{I}_{in}$ such that $P_{\mathcal{T}}$ is true in $\mathcal{I}$ and J justifies $P_{\mathcal{T}}$ over Δ_{gd}. Then there exists a model $\mathcal{M}$ of $\mathcal{T}$ that expands $\mathcal{I}|_S$ with S the set of defined literals in $J_{P_{\mathcal{T}}}$.*

Failure to find a model of Δ_{g} expanding $\mathcal{I}_{in}$ in which $P_{\mathcal{T}}$ is true implies the lack of models of $\mathcal{T}$ expanding $\mathcal{I}_{in}$. The former is the case when an inconsistency is detected in Δ_{g} after initializing $\mathcal{I}$ with $P_{\mathcal{T}}$. If Δ_{g} has no such models, then it has an unsatisfiable core, i.e., a set of rules from Δ_{g} such that no model exists that expands $\mathcal{I}$. Hence, it is also an unsatisfiable core for $\mathcal{T} = (P_{\mathcal{T}}, \Delta)$. To find an unsatisfiable core, one can, for example, use techniques described in (Torlak, Chang, & Jackson, 2008).

The above formalization is for a theory $\{P_{\mathcal{T}}, \Delta\}$ possibly containing an inductive definition. In the simpler case that the original input theory $\mathcal{T}$ was a set of FO sentences $\{\phi_1, \ldots, \phi_n\}$, the above method boils down to maintaining a grounded theory $\mathcal{T}_g$ and a non-ground theory $\mathcal{T}_d$, both consisting of FO sentences. For the non-ground part, a justification manager can maintain direct justifications as sets of literals that make the sentences of

$\mathcal{T}_g$ true. Major simplifications are that the justification manager never needs to find justifications that make a formula false and need not guard for cycles in the justification graph but only maintain that the direct justifications are consistent with one another and with $\mathcal{I}$.

Example 3.6. Consider the theory consisting of the sentences $P \Rightarrow Q$ and $Q \Rightarrow \forall x \in D : R(x)$. We can then keep both sentences in the non-grounded part by taking $\{\neg P\}$ as justification for the former sentence and $\{\neg Q\}$ for the latter. The union of both is consistent and hence would result in an acceptable state for the empty interpretation.

3.2 Practical constructions for FO(ID) theories

Roughly speaking, our lazy model expansion framework consists of two components. On the one hand, a standard model expansion algorithm that operates on $\{P_{\mathcal{T}}, \Delta_{\mathrm{g}}\}$ and, on the other hand, a justification manager that maintains a justification over Δ_{gd} and lazily grounds Δ_{d}. Lazy model expansion performs search over the space of acceptable states and aims at reaching a state where Theorem 3.4 (or Corollary 3.5) is applicable. To avoid slowing down the search during model expansion, the work done by the justification manager and the lazy grounding must be limited. To achieve this, we have designed a system in which the justification manager has no access to the grounded definition Δ_{g} and need not restore its state when the search algorithm backtracks over the current structure $\mathcal{I}$. The justification manager only has access to $\mathcal{I}$ and maintains justifications that are restricted to Δ_{d}. In particular, a literal defined in Δ_{g} is not allowed in a direct justification. Our justification manager maintains the following properties:

- Literals in direct justifications are either open in Δ_{gd} or defined in Δ_{d}.
- All direct justifications in J are kept consistent with each other and with the current structure $\mathcal{I}$.
- The justification graph defined by J has no positive or mixed cycles and is total.

To distinguish acceptable states that meet these additional requirements from acceptable states as defined in Definition 3.1, we call them *default acceptable states*; they can be defined as:

Definition 3.7 (Default acceptable state)**.** A state $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ is a default acceptable state if it is an acceptable state and, in addition, (**i**) literals in

direct justifications are either open in Δ_{gd} or defined in Δ_{d}, and (**ii**) J justifies the set of all literals for which J is defined.

It follows that default acceptable states satisfy two extra conditions: they do not justify literals defined in Δ_{d} in terms of literals defined in Δ_{g}, and the set of all literals in J is consistent. For an acceptable state it suffices that those in J_L are consistent. Since default acceptable states are acceptable states, Theorem 3.4 and Corollary 3.5 also hold for default acceptable states.

During standard model expansion, the main state-changing operations are expanding $\mathcal{I}$ by making literals true, either through choice or propagation, and by backjumping. When $S = \langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ is a default acceptable state and model expansion modifies $\mathcal{I}$ into $\mathcal{I}$', the new state $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I}' \rangle$ is not necessarily a default acceptable state. The following propositions identify situations where acceptability is preserved.

Proposition 3.8. *Let $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ be a default acceptable state, L a set of literals unknown in $\mathcal{I}$ and $\mathcal{I}'$ the consistent structure $\mathcal{I} \cup L$. If (**i**) literals of L either are not defined in Δ_d or have a direct justification in J and (**ii**) no direct justification in J contains the negation of a literal in L, then $\langle \Delta_g, \Delta_d, J, \mathcal{I}' \rangle$ is a default acceptable state.*

Proof. As the literals true in $\mathcal{I}$ and defined in Δ_{d} have a direct justification, it follows from (**i**) that all literals true in $\mathcal{I}$' and defined in Δ_{d} have a direct justification. As justifications in J is consistent with $\mathcal{I}$, then, by (**ii**), they are also consistent with $\mathcal{I}$'. Hence, J justifies all literals true in $\mathcal{I}$' and defined in Δ_{d}. □

Proposition 3.9. *Let $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ be a default acceptable state. Then $\langle \Delta_g, \Delta_d, J, \mathcal{I}' \rangle$ with $\mathcal{I}' <_p \mathcal{I}$ is a default acceptable state.*

Proof. The justification J justifies all literals defined in Δ and true in $\mathcal{I}$. As $\mathcal{I}$' is a subset of $\mathcal{I}$, J justifies all literals defined in Δ and true in $\mathcal{I}$'. □

The justification J has to be revised when model expansion adds a literal l to $\mathcal{I}$ and the conditions of Proposition 3.8 are not satisfied. The justification manager attempts to revise the direct justifications of literals where the negation of l occurs in their direct justification. When l is defined in Δ_{d} and has no direct justification, the justification manager attempts to find one. When the justification manager fails to revise J such that an acceptable state is obtained, the rule defining the offending literal is grounded and moved to Δ_{g}. If possible, the rule is first split up and only a small part is

grounded. At that point, propagation can interrupt the justification manager and infer the truth of additional literals, or detect an inconsistency, at which point the model expansion algorithm performs backjumping. In both cases, the justification manager has to resume the revision of the justification until an acceptable state is reached. Once an acceptable state is reached, model expansion can select and assign a choice literal. These processes are described in more detail in the next section.

As described above, we do not allow literals defined in Δ_{g} in the direct justifications of literals defined in Δ_{d}. The reason is that forbidden loops over both Δ_{g} and Δ_{d} can only be detected by also maintaining a justification in Δ_{g} (which we currently do not do). However, it is feasible to relax this condition by performing some static analysis over the rules of the definition. One case is when the body of a rule has no defined literals: literals defined by such a rule cannot be part of a cycle. A step further is to analyze the dependency graph: a literal defined in Δ_{g} is allowed in the direct justification of another literal when both literals do not belong to the same strongly connected component of the dependency graph. Indeed, in that case, they cannot be part of the same cycle.

3.3 An example

Our example uses a theory $\mathcal{T}$ which states that a symmetric graph ($edge/2$) exists where at least one node is reachable (predicate $R/1$) from the root node ($root$). Both $edge$ and $root$ are not interpreted, the input structure $\mathcal{I}$ only interprets the domain D as the set of domain elements $\{d_1, \ldots, d_n\}$.

The justification manager uses a *local* approach; initially, no literal has a direct justification. When the justification J does not meet the requirements for an acceptable state, a revision is made locally: the manager tries to find a direct justification for the offending literal without modifying direct justifications of other literals.

$P_{\mathcal{T}}$

$$\left\{\begin{array}{lll}
P_{\mathcal{T}} & \leftarrow C_1 \wedge C_2 & (1)\\
C_1 & \leftarrow \exists x \in D : \neg root(x) \wedge R(x) & (2)\\
C_2 & \leftarrow \forall (x\ y) \in D^2 : edge(x,y) \Rightarrow edge(y,x) & (3)\\
\forall x \in D : root(x) & \leftarrow x = d_1 & (4)\\
\forall x \in D : R(x) & \leftarrow root(x) \vee \exists y \in D : edge(x,y) \wedge R(y) & (5)
\end{array}\right\}$$

The lazy MX algorithm, sketched above, proceeds as follows:

1. The initial default acceptable state is $\langle \Delta_g, \Delta_d, J, \mathcal{I} \rangle$ in which Δ_g and $\mathcal{I}$ are empty, J is empty and $\Delta_d = \Delta$.

2. Propagation over $\{P_\mathcal{T}, \Delta_g\}$ sets $\mathcal{I}$ to $\{P_\mathcal{T}\}$. $P_\mathcal{T}$ is defined in Δ_d so the state becomes unacceptable as $P_\mathcal{T}$ is not justified in J. Setting the direct justification of $P_\mathcal{T}$ to $\{C_1, C_2\}$ does not help as the justification for $P_\mathcal{T}$ is not total in Δ_d; instead rule (1) is moved to Δ_g.

3. Unit propagation sets $\mathcal{I}$ to $\{P_\mathcal{T}, C_1, C_2\}$. Now C_1 and C_2 have to be justified. Consider first C_2 and rule 3. As *edge* is open, we can take a direct justification that sets all negative *edge* literals true (setting all positive *edge* literals true would be equally good). This justifies C_2 and avoids the grounding of the rule defining C_2.

4. Literal C_1 cannot be justified (with the local approach). However, as rule 2 is existentially quantified, one can avoid grounding the whole rule. One can perform a Tseitin transformation to isolate one instance and then only ground that instance. For the purpose of illustration, we make the worst choice possible and instantiate x with d_1:

$$\left\{ \begin{array}{lll} C_1 & \leftarrow (\neg root(d_1) \wedge R(d_1)) \vee T & (2a) \\ T & \leftarrow \exists x \in D \setminus \{d_1\} : \neg root(x) \wedge R(x) & (2b) \end{array} \right\}$$

   Rule 2a is moved to Δ_g and a default acceptable state is reached.

5. No further propagation is possible, so a choice has to be made. As C_1 is true, the body of rule 2a has to become true. Preferably not selecting a Tseitin (this would trigger more grounding), the first disjunct is selected by model expansion and propagation extends the structure with $root(d_1) = \mathbf{f}$ and $R(d_1) = \mathbf{t}$. The literal $root(d_1)$ is defined in Δ_d by rule 4 but cannot be justified with the local approach. To avoid fully grounding it, it is transformed in rule 4a that defines $root(d_1)$ and rule 4b that defines *root* for the other domain elements:

$$\left\{ \begin{array}{lll} root(d_1) & \leftarrow \top & (4a) \\ \forall x \in D \setminus \{d_1\} : root(x) & \leftarrow x = d_1 & (4b) \end{array} \right\}$$

   Rule 4a is moved to Δ_g. Note that it has no defined literals in the body, hence it is safe to use $root(d_1)$ in direct justifications. As a consequence of extending Δ_g, propagation detects an inconsistency. After backtracking to $\mathcal{I} = \{P_\mathcal{T}, C_1, C_2\}$, the subsequent propagation sets the structure $\mathcal{I}$ to $\{P_\mathcal{T}, C_1, C_2, root(d_1), T\}$. Still not in a default

acceptable state (T is not justifiable using the local approach), rule 2b is further transformed to split off another instance.

$$\left\{ \begin{array}{lll} T & \leftarrow (\neg root(d_2) \wedge R(d_2)) \vee T_2 & (2ba) \\ T_2 & \leftarrow \exists x \in D \setminus \{d_1, d_2\} : \neg root(x) \wedge R(x) & (2bb) \end{array} \right\}$$

Rule 2ba is moved to Δ_{g}, while rule 2bb remains in Δ_{d}.

6. Again, the search avoids the new Tseitin, selecting the first disjunct in rule 2ba which propagates $\neg root(d_2)$ and $R(d_2)$. The literal $\neg root(d_2)$ is defined in Δ_{d}, but can be justified by $\top$ (the body of rule 4b is false for $x = d_2$). The literal $R(d_2)$ cannot be justified (as all *edge* literals are false in the current justification graph) and rule 5 is transformed to split off the d_2 instance. Actually, this instance in turn has a disjunctive body with a complex subformula, so to avoid grounding the subformula, we break it up in two parts and introduce another Tseitin.

$$\left\{ \begin{array}{lll} R(d_2) & \leftarrow root(d_2) \vee T_3 & (5aa) \\ T_3 & \leftarrow \exists y \in D : edge(d_2, y) \wedge R(y) & (5ab) \\ \forall x \in D \setminus \{d_2\} : R(x) & \leftarrow root(x) & \\ & \quad \vee \exists y \in D : edge(x, y) \wedge R(y) & (5b) \end{array} \right\}$$

Rule 5aa is moved to Δ_{g}, the others remain in Δ_{d}.

7. The current structure $\mathcal{I}$ is $\{P_{\mathcal{T}}, C_1, C_2, root(d_1), T, \neg root(d_2), R(d_2)\}$, hence propagation on rule 5aa in Δ_{g} extends it with T_3. However, T_3 cannot be justified and rule 5ab is transformed to avoid its full grounding. We split off the d_1 case

$$\left\{ \begin{array}{lll} T_3 & \leftarrow (edge(d_2, d_1) \wedge R(d_1)) \vee T_4 & (5aba) \\ T_4 & \leftarrow \exists y \in D \setminus \{d_1\} : edge(d_2, y) \wedge R(y) & (5abb) \end{array} \right\}$$

Rule 5aba is moved to Δ_{g} while rule 5abb remains in Δ_{d}.

8. Again search selects the first disjunct and propagates T_3, $edge(d_2, d_1)$ and $R(d_1)$. The literal $R(d_1)$ is defined in Δ_{d}, but $root(d_1)$ is a direct justification for it. While $root(d_1)$ is defined in Δ_{g}, it is safe to use it in a direct justification as discussed earlier. The literal $edge(d_2, d_1)$ is in conflict with the direct justification for C_2 (rule 3). To handle this conflict, we split off the affected instance ($x = d_2, y = d_1$). Literal T_5 inherits the unaffected part of the direct justification of C_2 (all negative literals over *edge* except $edge(d_2, d_1)$).

$$\left\{ \begin{array}{lll} C_2 & \leftarrow (edge(d_2, d_1) \Rightarrow edge(d_1, d_2)) \wedge T_5 & (3a) \\ T_5 & \leftarrow \forall (x\ y) \in D^2 \setminus \{(d_2, d_1)\} : edge(x, y) \Rightarrow edge(y, x) & (3b) \end{array} \right\}$$

Rule 3a is moved to Δ_{g} while rule 3b remains in Δ_{d}.

9. Propagation on rule 3a extends $\mathcal{I}$ with $edge(d_1, d_2)$ and T_5. The literal $edge(d_1, d_2)$ is in conflict with the direct justification for T_5 (rule 3b) and part of rule 3b has to be grounded. First, we split off the instance $\{x = d_1, y = d_2\}$ as follows.

$$\left\{\begin{array}{lll} T_5 & \leftarrow (edge(d_1, d_2) \Rightarrow edge(d_2, d_1)) \wedge T_6 & (3ba) \\ T_6 & \leftarrow \forall (x\ y) \in D^2 \setminus \{(d_2, d_1), (d_1, d_2)\} : & \\ & \qquad edge(x, y) \Rightarrow edge(y, x) & (3bb) \end{array}\right\}$$

   Rule 3ba is moved to Δ_{g} while rule 3bb remains in Δ_{d} and inherits the remainder of the direct justification (all negative literals over *edge* except $edge(d_1, d_2)$ and $edge(d_2, d_1)$). Propagation on rule 3ba extends $\mathcal{I}$ with T_6; while defined in Δ_{d}, it is justified.

By now, Δ_{g} consists of the rules 1, 2a, 4a, 2ba, 5aa, 5aba, 3a, and 3ba, and the residual definition Δ_{d} consists of the rules 4b, 2bb, 5b, 5abb, and 3bb. The current structure $\mathcal{I}$ is $\{P_{\mathcal{T}}, C_1, C_2, root(d_1), \neg root(d_2), edge(d_2, d_1), edge(d_1, d_2), R(d_1), R(d_2), T, T_3, T_5, T_6\}$ and is a model of $P_{\mathcal{T}} \cup \Delta_{\mathrm{g}}$.

Of these literals, $\neg root(d_2)$, $R(d_1)$ and T_6 are defined in Δ_{d}. Literal $\neg root(d_2)$, defined by rule 4b has $\{\top\}$ as direct justification. Literal $R(d_1)$, defined by rule 5b, has $\{root(d1)\}$ as direct justification. Literal T_6, defined by rule 3bb has as direct justification the set of all negative *edge* literals over D except $edge(d_1, d_2)$ and $edge(d_2, d_1)$. To obtain a full model of the theory, $\mathcal{I}$ is extended with the literals of the above direct justifications. In this case, this assigns all open literals and the model can be completed by the well-founded model computation over Δ_{gd}. Actually, this can be done without grounding the definition (Jansen, Jorissen, & Janssens, 2013).

4 Justification management

In Section 3.2, we have instantiated our general framework, developed in Section 3.1, for a justification manager that only has access to Δ_{d}. In the example of Section 3.3, the justification was constructed on demand, i.e., each time some literal needed a (different) direct justification, the body of its defining rule was analyzed and a justification was extracted. If that failed, part of the rule was grounded. This was called the *local approach*. One can also imagine a *global approach*, where more rules of Δ_{d} are considered at once in an attempt to select direct justifications that minimize the grounding of the rules as a whole. Obviously, a global approach will be more

time consuming, so should not be applied every time an adjustment of the justification is required. In this section, we describe both approaches.

Before describing the algorithms, we introduce some notations and assume some normalizations have been done. Quantifiers range over a single variable. Variable names are not reused in a formula. As stated in the preliminaries, literals $\top$ and $\bot$ are interpreted as respectively $\mathbf{t}$ and $\mathbf{f}$. We assume that some basic reductions are always applied on formulas, e.g., $\bot \wedge \varphi$ reduces to φ, $\forall x \in \emptyset : \varphi$ reduces to $\mathbf{t}$, ... The function nnf reduces a formula to its negation normal form. With S a set and s a single element, $S + s$ and $S - s$ are used as shorthands for $S \cup \{s\}$ and $S \backslash \{s\}$.

4.1 The local approach

Algorithm 2 shows the top level of the lazy_mx model expansion algorithm. The theory passed to the algorithm consists of the fact $P_{\mathcal{T}}$ and the definition Δ; $\mathcal{I}_{in}$ is the input structure to be expanded. Definitions Δ_{d} and Δ_{g} are initialized with Δ and the empty definition, respectively, and $\mathcal{I}$ is initialized with $\mathcal{I}_{in}$. The set of ground sentences T_{g} is initialized with the fact $P_{\mathcal{T}}$ and the initial justification J is empty. An auxiliary (FIFO) queue q_{ch} is initialized as empty. This queue keeps track of literals for which the direct justification needs to be checked.

The main loop performs model expansion over $T_{\mathrm{g}} \cup \Delta_{\mathrm{g}}$, interleaved with work by the justification manager towards establishing a default acceptable state. The model expansion part consists of propagation (the call to propagate), the test for inconsistency of the current state (with clause learning and backjumping), the test whether a model of $T_{\mathrm{g}} \cup \Delta_{\mathrm{g}}$ is found (returning the model and the justification) and of a choice step that selects an unknown literal in $T_{\mathrm{g}} \cup \Delta_{\mathrm{g}}$ and assigns it a value. The test for a model is only performed in a default acceptable state (i.e., when the queue q_{ch} is empty) as this ensures the well-founded model computation can expand the current structure $\mathcal{I}$ —extended with the direct justifications of all literals— into a model of the whole theory. Also the choice step only takes place in a default acceptable state, to ensure that the search space is limited to the state space of default acceptable states. Literals that become assigned (by propagation or choice) need a direct justification if they are defined in Δ_{d}, hence they are queued by `checkliteral` if they do not have one. Also, assigned literals of which the negation is used in a direct justification need to be queued, as this direct justification is now invalid. When no propagation is possible and the current structure is consistent, a literal is removed from the queue and its direct justification is updated by the lazy_ground function. Literals are

Algorithm 2: The lazy_mx lazy model expansion algorithm.

```
1  Function lazy_mx (atomic sentence P_T, definition Δ, structure I_in)
       Output: either a model of Δ_g and J or false
2      T_g := {P_T }; Δ_g := ∅; Δ_d := Δ; J := ∅; I := I_in; q_ch := ∅;
3      while true do
4          L := propagate(T_g ∪ Δ_g, I);
5          I := I ∪ L;
6          foreach l ∈ L do q_ch:=check_literal(l,q_ch);
7          if I is inconsistent then
8              T_g += learn nogood;
9              if conflict at root level then return false;
10             I := I at state of backjump point;
11         else if q_ch is not empty then
12             (l, q_ch) := dequeue(q_ch);
13             lazy_ground(l);
14         else if I is a model of T_g ∪ Δ_g then
15             return I, J;
16         else
17             select choice literal l; I := I + l;
18             q_ch:=check_literal(l,q_ch);
19 Function checkliteral(literal l, literal queue q_ch)
       Data: global Δ_d and J    Output: updated queue
20     if l defined in Δ_d and J(l) = undef then  q_ch := enqueue(l,q_ch) ;
21     foreach l' such that ¬l ∈ J(l') do  q_ch := enqueue(l',q_ch) ;
22     return q_ch;
```

processed one at a time as lazy_ground may extend Δ_g, which may allow for extra propagation. As propagation is a fast and powerful operation, it is given priority over updating the justifications. Second, propagation might result in conflict and backtracking, after which some justifications might not need an update anymore, as discussed below.

4.1.1 Lazy grounding of one rule

The function lazy_ground, Algorithm 3, checks the direct justification of its input literal. As the lazy model expansion algorithm may have backtracked since the literal was queued, it may well be that the direct justification of the literal is again acceptable as it is. Otherwise, namely when the literal is true in $\mathcal{I}$, defined in Δ_d and has no direct justification, the function justify, Algorithm 4, is called to construct a new direct justification. If a direct justification already existed, justify first checks whether it is still valid. If justify did not find a valid direct justification, part of the rule has to be grounded, which is done through the call to split_and_ground. The old justification, if any, is passed as argument to split_and_ground, as we show that it can help to guide the grounding process.

Algorithm 3: The lazy grounding of a literal.

```
1  Function lazy_ground (literal l)
       Data: global Δ_d, J and I
2      success := true;
3      j_old := J(l);
4      if l ∈ I and l defined in Δ_d and J(l) = undef then
5          success := justify(l);
6      else if J(l) ≠ undef then
7          success := justify(l);
8      if success = false then
9          J := J[l → false];
10         split_and_ground(l, j_old);
```

Before going into more details, we first analyze which properties we want to maintain in the justification. The direct justifications of literals in the q_{ch} queue might be invalid, hence they should not be considered as part of the current justification. The global invariants we maintain are:

- neither positive nor mixed cycles (negative ones are allowed),

- literals in the justification are consistent with each other.

For direct justifications of literals not on the queue, we maintain that

- they do not contain literals defined in Δ_{g} (unless such a literal is safe, i.e., it can never be part of a cycle),
- the literals they contain that are defined in Δ_{d} either have a direct justification or are on the queue themselves.

These invariants imply that a default acceptable state is reached when the q_{ch} queue is empty. Indeed, it follows from the invariants that the current justification is total in that situation and hence all literals that have a direct justification are justified (Definition 2.6). Due to the policy followed to queue literals, the current justification is also consistent with $\mathcal{I}$ while all literals true in $\mathcal{I}$ and defined in Δ_{d} have a justification, hence $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$ is a default acceptable state.

Algorithm 4: The justify algorithm.

```
1 Function justify (literal l)
      Data: global Δ_d and J    Result: update to J
      Output: true if a justification was found, false otherwise
2     if J(l) exists and obeys the invariants then return true ;
3     φ := body of the rule defining l;
4     if l is a negative literal then φ := nnf(¬φ) ;
5     dj := build_djust(l, φ, init_just(l));
6     if dj = false then return false ;
7     else J:= J[l → dj]; return true ;
```

The function justify, described in Algorithm 4 starts by checking whether the direct justification of the literal obeys the invariants on direct justifications and on the justification; if so, nothing needs to be done. Otherwise, build_djust is called, a function that attempts to find a valid direct justification. If found, the justification is updated and *true* is returned; if not, *false* is returned.

4.1.2 Building a direct justification

The purpose of build_djust, Algorithm 5, is to find a direct justification for literal l that obeys the invariants of the direct justifications as well as the global invariants on the justification graph extended with this direct

justification. It is a recursive function which takes three parameters: (**i**) the literal l, (**ii**) the formula to be made true by the direct justification (initially the whole body of the rule defining the literal; note that the initialization takes the negation of the rule when the literal is negative), (**iii**) a description of the direct justification derived so far, initialized through $\mathsf{init_just}(l)$.

Algorithm 5: The build_djust algorithm.

```
1  Function build_djust (literal l, formula φ and justification ⟨L, B⟩)
       Input: B binds all free variables of φ
       Output: either a direct justification or false
2      switch φ do
3          case φ is a literal
4              if valid(l, ⟨L ∪ {φ}, B⟩) then return ⟨L ∪ {φ}, B⟩;
5              else return false;
6          case ∀x ∈ D : ψ
7              return build_djust(l, ψ, ⟨L, B ∪ {x ∈ D}⟩);
8          case ∃x ∈ D : ψ
9              foreach d_i ∈ D do
10                 ⟨L′, B′⟩ := build_djust(l, ψ, ⟨L, B ∪ {x ∈ d_i}⟩);
11                 if ⟨L′, B′⟩ ≠ false then return ⟨L′, B′⟩;
12             return false;
13         case φ_1 ∧ ... ∧ φ_n
14             foreach i ∈ [1, n] do
15                 ⟨L′, B′⟩ := build_djust(l, φ_i, ⟨L, B⟩);
16                 if ⟨L′, B′⟩ = false then return false;
17                 else ⟨L, B⟩ := ⟨L′, B′⟩;
18             return ⟨L, B⟩;
19         case φ_1 ∨ ... ∨ φ_n
20             foreach i ∈ [1, n] do
21                 ⟨L′, B′⟩ := build_djust(l, φ_i, ⟨L, B⟩);
22                 if ⟨L′, B′⟩ ≠ false then return ⟨L′, B′⟩;
23             return false;
```

Before going into detail, we discuss how to represent direct justifications. Basically, we could represent a direct justification as a set of ground literals. However, this set can be quite large and using a ground representation might hence defy the purpose of lazy grounding. Instead, we represent a direct

justification as a pair $\langle L, B\rangle$ with L a set of possibly non-ground literals and B a set of bindings $x_i \in D_i$ with x_i a variable and D_i a domain. A set of bindings $\{x_1 \in D_1, \ldots, x_n \in D_n\}$ represents the set of variable substitutions $S_B = \{\{x_1/d_1, \ldots, x_n/d_n\} \mid d_i \in D_i \text{ for each } i \in [1, n]\}$. The set of ground literals represented by $\langle L, B\rangle$ is then $\{l\theta \mid l \in L \text{ and } \theta \in S_B\}$. The direct justification of a literal $P(\overline{d})$, defined by a rule $\forall \overline{x} \in \overline{D} : P(\overline{x}) \leftarrow \varphi$, is initialized by $\mathsf{init_just}(l)$ as $\langle \emptyset, \{x_1 \in \{d_1\}, \ldots, x_n \in \{d_n\}\}\rangle$. In effect, B selects the appropriate variable instantiation from the domains to identify the relevant rule instantiation, while the set of literals is empty.

The build_djust algorithm searches for a set of literals making φ true. It works by recursively calling itself on subformulas of φ and composing the results afterwards into a larger justification for φ.

The base case is when the formula is a literal. To make that literal true, all instances of the literal under binding B must be true, hence, the set of literals L is extended with the literal itself. The resulting direct justification has to satisfy all invariants, which is checked by the call to valid: it returns *true* for a call $\mathsf{valid}(l, dj)$ if dj satisfies the invariants to be (part of) a direct justification for l and $J[l \to dj]$ satisfies the invariants on the justification.

A universally quantified formula $\forall x \in D : \psi$ has to be true for each instance of the quantified variable. Hence, in the recursive call, the binding B is extended with $x \in D$. For an existentially quantified formula, it suffices that one instance is true. Hence, a minimal approach is to try each instance separately until one succeeds; if all fail, *false* is returned. Note however that we do not want to iterate over each domain element if D is large, which would be similar to constructing the grounding itself. Instead, if D is large, we extend the binding with $x \in D$.

Conjunction is similar to universal quantification, except that explicit iteration over each conjunct is needed. As soon as one conjunct fails, the whole conjunction fails. Disjunction is similar to existential quantification.

Note that the order in which build_djust iterates over subformulas and instantiations has an effect on the justification that is found.

Example 4.1. Consider the following rule over a large domain D.

$$H \leftarrow \forall x \in D : \neg P(x) \vee (\exists y \in D : Q(x, y) \wedge \neg R(x, y))$$

Assume that J is empty and we have no loops to keep track of. Applying build_djust to H then results in either the justification $\{\neg P(x) \mid x \in D\}$ or the justification $\{Q(x, y) \mid x \in D, y \in D\} \cup \{\neg R(x, y) \mid x \in D, y \in D\}$, depending on the order of iteration.

4.1.3 Partially grounding a rule

The last bit of the lazy model expansion algorithm handles the case where no (new) justification could be found for a defined literal l. A straightforward solution would be to call one_step_ground on the associated rule, add the resulting ground rule(s) to Δ_{g} and add rules over newly introduced (Tseitin) symbols to Δ_{d}. Indeed, these new symbols are unknown in $\mathcal{I}$, so can be kept safely in Δ_{d} for now, even without a direct justification. However, in many cases such an operation results in too much grounding.

Example 4.2. Consider a rule r_1 of the form $\forall x \in D : P(x) \leftarrow \varphi$ in a situation where no justification can be found for atom $P(d)$. Straightforwardly applying one_step_ground to r_1 would instantiate x with all elements in D, resulting in $|D|$ rules, while in fact it would suffice to split r_1 in two rules, one for the instance $x = d$ and one for the remainder. Another example applies to a rule r_2 of the form $H \leftarrow \forall x \in D : Q(x) \vee R(x)$ and a direct justification $J[H] = \langle Q(x), \{x \in D\}\rangle$. When $Q(d)$ becomes false (assume R is not a valid alternative justification), one_step_ground instantiates x with all elements in D. Instead, it is be better to split off the instantiation $x = d$ and introduce a Tseitin for the remainder. For the latter, removing $Q(d)$ from the original justification results in a valid direct justification.

The split_and_ground algorithm (Algorithm 6) has to ground part of the rule defining a given literal l, say $P(\overline{d})$. The first step is to split off the rule instance for which the rule defining l has to be grounded (the call to split). Let $\forall \overline{x} \in \overline{D} : P(\overline{x}) \leftarrow \varphi$ be the rule in Δ_{d} that defines $P(\overline{d})$. We then replace that rule by $\forall \overline{x} \in \overline{D} - \overline{d} : P(\overline{x}) \leftarrow \varphi$ in Δ_{d} and additionally return the rule $P(\overline{d}) \leftarrow \varphi[\overline{x}/\overline{d}]$. Afterwards, we apply one_step_ground and add the rules to their corresponding definitions.[1]

Algorithm 6: The split_and_ground algorithm.

```
1 Function split_and_ground (literal l, justification j_old)
      Input: l is defined in Δ_d, j_old is the previous justification (if any)
      Result: update to Δ_g, Δ_d, J, and q_ch
2     r := split(l, j_old);
3     (Δ'_g, Δ'_d) := one_step_ground(r);
4     Δ_g ∪= Δ_g'; Δ_d ∪= Δ_d';
```

[1] Recall, rules returned by one_step_ground always have a unique head. Those in Δ_{d}' even a unique new Tseitin as head.

The result of split_and_ground is that definition Δ_{gd} that is "more" ground than the previous one. The limit is a ground definition Δ_{gd} in which Δ_{d} is empty and Δ_{g} is strongly $voc(\Delta)$-equivalent with Δ.

Even if no justification was found, we can do better than just splitting off l and applying one_step_ground, as shown in Example 4.2. First, splitting can be made significantly more intelligent, which is discussed in Section 4.1.5. Second, we can improve one_step_ground to only ground part of expressions if possible, which we described below.

Improving one_step_ground Applying one_step_ground to a rule $l \leftarrow \varphi$ iterates over all subformulas/instantiations of φ. For example if φ is the sentence $\exists x \in D : P(x)$, the result are $|D|$ new rules and as many new Tseitin symbols. Instead, depending on the value of l, we can suffice with only introducing one (or some) of these subformulas, as shown in Algorithm 7, which extends the switch statement of one_step_ground with two higher-priority cases. If l is true, we can suffice by grounding one disjunct/existential instantiation and delay the rest by Tseitin introduction. If l is false, we take a similar approach for conjunction/universal quantification.

Algorithm 7: Additional cases for the one_step_ground algorithm.

```
switch r do
    case l ← φ_1 ∨ ... ∨ φ_n and I(l) ≠ f
        choose i ∈ [1, n];
        Δ_d += T ← ⋁_{j∈{1...n}−i} φ_j ;
        r := L ← φ_i ∨ T;
    case l ← ∃x ∈ D : φ and I(l) ≠ f
        choose d ∈ D;
        Δ_d += T ← ∃x ∈ D\d : φ;
        r := l ← φ[x/d] ∨ T;
    analogous cases for ∧ and ∀ and I(l) ≠ t
```

4.1.4 Algorithmic properties

Correctness and termination of the presented algorithms is discussed in the following theorem.

Theorem 4.3 (Correctness and termination)**.** *If algorithm lazy_mx returns an interpretation $\mathcal{I}$, then expanding $\mathcal{I}$ with J, applying the well-founded*

evaluation over Δ_{gd} and restricting it to $voc(\mathcal{T})$ results in a model of $\mathcal{T}$. If the algorithm returns false, no interpretation exists which is more precise than $\mathcal{I}_{in}$ and satisfies $\mathcal{T}$.

Algorithm lazy_mx *terminates if $\mathcal{I}$ is finite. Otherwise, termination is possible but not guaranteed.*[2]

Proof. If lazy_mx returns an interpretation $\mathcal{I}$, $\mathcal{I}$ is a model of Δ_{g} and q_{ch} is empty. Given the properties of split_and_ground, after applying lazy_ground for a literal l, no more constructions are violated by l and any rule defining l has a construction valid in $\mathcal{I}$ or is part of Δ_{g}. Hence if q_{ch} is empty, we are in a default acceptable state. In that case, the model construction is correct as shown in 3.4. If lazy_mx returns *false*, it has been proven that Δ_{g} has no models in $\mathcal{I}_{in}$. In that case, there can also be no models of Δ_{gd} and hence $\mathcal{T}$ also has no models expanding $\mathcal{I}_{in}$.

Without calls to lazy_ground, the search algorithm terminates for any finite Δ_{g}. lazy_ground itself produces an ever-increasing ground theory Δ_{g} with the full grounding as limit. Hence, lazy_mx always terminates if $\mathcal{I}_{in}$ is finite. If $\mathcal{I}_{in}$ is infinite, the limit of Δ_{g} is an infinite grounding, so termination cannot be guaranteed. □

4.1.5 Symbolic justifications, incremental querying and splitting

The algorithms presented above are sound and complete. There are however two important observations that create room for improvements. In both cases, it revolves around not treating justifications as just a set of literals but taking the formula from which they were derived into account.

Incremental querying First, we observe that if multiple justifications exist for (subformulas of) a formula φ, we can delay grounding even more.

Example 4.4. Consider the formula $\forall x \in D : P(x) \vee Q(x)$, where both $\langle\{P(x)\}, \{x \in D\}\rangle$ and $\langle\{Q(x)\}, \{x \in D\}\rangle$ are justifications. From that, we could derive the justification: for each $d \in D$, make either $P(d)$ *or* $Q(d)$ true. Hence, *only* when for one d, both $P(d)$ and $Q(d)$ become false is more grounding necessary.

We can do this automatically by changing build_djust as follows.

[2] It is possible to change the integration of lazy_ground in lazy_mx to guarantee termination if a finite model exists, see Section 5.2.

- The algorithm is allowed to select multiple disjunctions / existential quantifications even if a valid justification was already found for one (Lines 11 and 22).
- build_djust does not build a justification, but builds a *justification formula* from the subformulas/instantiations it selects. For example, in the $\forall$-case, **return** build_djust$(l, \psi, \langle L, B \cup \{x \in D\}\rangle)$ is changed to **return** $\forall x \in D :$ build_djust$(l, \psi, \langle L, B \cup \{x \in D\}\rangle)$.
- The validity check (valid) is extended to only return true if the justification formula is not false.

It is straightforward to see that a justification formula ψ of a formula φ entails φ. From ψ, the justification itself can be derived directly as the set of non-false literals in the (full) grounding of ψ.

Naturally, by allowing more complex formulas (instead of a conjunction of universally quantified literals), it becomes more expensive to track whether a formula has become false after changes to $\mathcal{I}$. This is in fact the *incremental query* problem. In the experiments, we limit the depth of the allowed formulas and use a straightforward (expensive) incremental query algorithm to track whether justification formulas have become false.

Body splitting Second, if a literal l violates the justification j of a formula φ, it is often possible to split j and φ each in two parts: a small part to which the violation really applies and a large remainder for which we can directly reuse (a reduced version of) j.

Example 4.5. Consider the rule $h \leftarrow \forall x \in D : P(x)$ with justification $\langle\{P(x)\}, \{x \in D\}\rangle$ for h true in $\mathcal{I}$. When $P(d)$ becomes true, it is easy to see that we can split off the "violating instantiation" by rewriting the original rule into $h \leftarrow P(d) \wedge T$ and adding the rule $T \leftarrow x \in D \backslash d : P(x)$. Crucially, a justification for the latter can be derived from the original justification, namely $\langle\{P(x)\}, \{x \in D \backslash d\}\rangle$. The latter rule can hence be added to Δ_{d} and its justification to J and we only are only left with the former rule.

This is the responsibility of the split operation, which takes as input a literal l defined in Δ_{d} by a rule $r = h \leftarrow \varphi$ and its old justification formula j_{old}. Next to splitting of the rule instance based on l (as described earlier), its task is to identify subformulas of φ for which $j_{old} \backslash l$ is a justification formula and apply Tseitin transformation on them.

Proposition 4.6. *Given a literal l in a justification j of a sentence φ. If the full grounding of a subformula ψ of φ does not contain $\neg l$, then $j \backslash l$ is a justification of ψ.*

Assume v is the domain literal that violated a justification and resulting in the current call to lazy grounding (v can be computed from the state or, better still, can be passed down from lazy_mx). Assume we know all variable instantiations $\{\overline{x} = \overline{d}_1, \ldots, \overline{x} = \overline{d}_n\}$ under which φ contains occurrences of a literal $\neg v$. Then the following algorithm can be used to efficiently split φ for v. For each variable instantiation, say $\overline{x} = \overline{d}_i$, visit r recursively and depth-first. Whenever a quantification $\forall x \in D : \varphi$ is encountered with x equal to $x_j \in \overline{x}$, replace it by $(\forall x \in D - d_j : \varphi) \wedge \varphi[x = d_j]$. Tseitin transformation is applied to the left-hand conjunct. This results in a new rule r_v for which $j_{old} \backslash v$ is a justification. Afterwards, the right-hand conjunct is visited recursively. The result is a set of new rules for which no new justification has to be sought and a smaller rule r' which are then passed to one_step_ground. Correctness follows from the fact the $j_{old} \backslash v$ is a valid justification, none of the new rules contains v, and from the correctness of Tseitin transformation.

Example 4.7. In Example 4.1, justifications were sought for H in the rule

$$H \leftarrow \forall x \in D : \neg P(x) \vee (\exists y \in D : Q(x, y) \wedge \neg R(x, y)).$$

Assume we selected the justification $j = \{Q(x, y) \mid x \in D, y \in D\} \cup \{\neg R(x, y) \mid x \in D, y \in D\}$. When $l = Q(d_1, d_2)$ becomes false, j is no longer consistent with $\mathcal{I}$; $j \backslash l$, however, is still consistent with $\mathcal{I}$, but is not a justification of the whole body. On the other hand, $j \backslash l$ is a justification for the subformula $\neg P(x) \vee \exists y \in D : Q(x, y) \wedge \neg R(x, y)$ for each instantiation of x different from d_1. Consequently, we can split the quantification $\forall x \in D :$ into $x \in D - d_1$ and $x = d_1$ and apply Tseitin transformation to the former. Afterwards, we recursively visit the latter formula and apply a similar reasoning to the existential quantification. The operations on the formula are illustrated in Figure 2. The result are the following rules, where the rule for H is now even ground.

$$\left\{ \begin{array}{l} H \leftarrow T_1 \wedge (\neg P(d_1) \vee (T_2 \vee (Q(d_1, d_2) \wedge \neg R(d_1, d_2)))) \\ T_1 \leftarrow \forall x \in D \backslash \{d_1\} : \neg P(x) \vee \exists y \in D : Q(x, y) \wedge \neg R(x, y) \\ T_2 \leftarrow \exists y \in D \backslash \{d_2\} : Q(d_1, y) \wedge \neg R(d_1, y) \end{array} \right\}$$

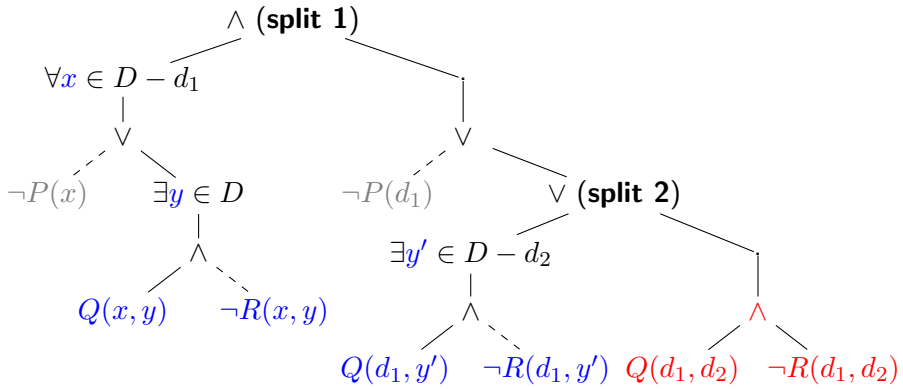

Figure 2: The rule $\forall x \in D : \neg P(x) \vee \exists y \in D : Q(x, y) \wedge \neg R(x, y)$ is split for a violating literal $Q(d_1, d_2)$. The original justification without $Q(d_1, d_2)$ is a justification of the left-hand side of all splits, with the justification formula shown in blue. The remaining non-justified formula is shown in red.

Obtaining all variable instantiations, required to apply split, can be done during an initial visit of φ by verifying all possible ways to obtain l. However, this information can be directly extracted from build_djust: for each selected literal occurrence l, we keep track of the path taken through the parse tree of φ: which subformulas/instantiations were selected to get to l.

Example 4.8. Assume the rule $C_1 \leftarrow \forall(x\ y) \in D^2 : \neg edge(x, y) \vee edge(y, x)$ has to be constructed true, J is empty and $\mathcal{I}$ does not interpret $edge$. If build_djust recursively visits the body of the rule until $\neg edge(x, y)$, $\neg edge(x, y)$ is returned as it is a valid literal to use. Going up one level, we store that for $\neg edge(x, y) \vee edge(y, x))$, we selected $\{\neg edge(x, y)\}$. Assuming no more disjuncts are selected, $\neg edge(x, y)$ is returned again. Going back up through both quantifications, we store that, for both quantifications, we selected D as the set of relevant domain elements, and build_djust returns the justification formula $\forall(x\ y) \in D^2 : \neg edge(x, y)$.

4.2 The global approach

Finding justifications using the greedy local approach can easily lead to more grounding than necessary. Consider for example the sentences $\forall x \in D : P(x)$ and $\forall x \in D : P(x) \Rightarrow Q(x)$. Applying the local approach to the second sentence first (with an empty J), might result in the construction which makes atoms over P false. Applying the local approach to the first sentence

then finds no valid justification, requiring to fully ground it. The global approach takes a *set* of rules as input and tries to select justifications for them such that the *expected* grounding size of the whole set is minimal.

We cast the task, called the *optimal justification* problem, as a problem on a graph as follows. The graph consists of two types of nodes R ("rule") and J ("justification"). Each rule node corresponds to a rule in Δ_{d} and a truth value (true, false or unknown) which indicates whether the head is true or false or whether it is kept in Δ_{d} but not be assigned a justification. Each justification node corresponds to a (symbolic) set of literals. Next, there are three types of edges V ("valid"), C ("conflict") and D ("depends on"). There is a valid edge from a justification node to true or false rule nodes that are directly justified by it. There is a conflict edge between rule nodes of the same rule and between justification nodes that contain opposite literals. There is also a conflict edge between a true (false) rule node that defines l and a justification node that contains $\neg l$ (l) and between an unknown rule node that defines l and a justification node that contains either l or $\neg l$. There is a depends edge from a true (false) rule node to a justification node if the justification contains negative (positive) literals defined by the rule. The idea is then to select rule and justification nodes such that

- Each selected true or false rule node is connected with a V edge to at least one selected justification (justified). Each selected justification is connected through a V edge to a selected true or false rule node.
- Selected nodes are not connected by a C edge (no conflicts).
- There are neither positive nor mixed cycles in the subgraph consisting of the D and V edges connecting selected nodes (the selected justifications justify the selected rules).

From a selection satisfying these constraints, an initial J can be built as follows. $J[(\neg)l \to j]$ if l is defined in r, its true (false) rule node was selected and j is the union of the justifications of its associated selected justification nodes. All literals defined by rules for which no rule node was selected are added to the initial q_{ch} queue, to be handled by the local approach.

This type of problem is somewhat related to the NP-hard *hitting set* (or *set cover*) problem (Karp, 1972): given a set of "top" and "bottom" nodes and edges between them, the task is to find a minimal set of bottom nodes such that each top node has an edge to at least one selected bottom node.

Given a default acceptable state $\langle \Delta_{\mathrm{g}}, \Delta_{\mathrm{d}}, J, \mathcal{I} \rangle$, the input for the optimal justification problem is generated as follows. For any rule in Δ_{d}, a node

is constructed for each of the three truth values (only one when the truth value is known in $\mathcal{I}$) and also their conflict edges are added. Justification nodes and their V edges are obtained using a (straightforward) adaptation of build_djust which returns a set of possible justifications that make the head of a rule true (false). E.g., for a rule $\forall \overline{x} \in \overline{D} : P(x) \leftarrow \varphi$, build_djust is called with the literal $P(x)$ and the binding is initialized at $\{\overline{x} \in \overline{D}\}$. Conflict and depends-on edges are derived by checking dependencies between justifications and between rules and justifications. To keep this efficient, it is done on the symbolic level.

Example 4.9. Consider the theory of our running example, after $P_{\mathcal{T}}$, C_1 and C_2 have been propagated to be true. Definition Δ_{d} is then

$$\left\{\begin{array}{lll} C_1 & \leftarrow \exists x \in D : \neg root(x) \wedge R(x) & (2) \\ C_2 & \leftarrow \forall (x\ y) \in D^2 : edge(x,y) \Rightarrow edge(y,x) & (3) \\ \forall x \in D : root(x) & \leftarrow x = d_1 & (4) \\ \forall x \in D : R(x) & \leftarrow root(x) \vee \exists y \in D : edge(x,y) \wedge R(y) & (5) \end{array}\right\}$$

The associated optimal construction set input is shown in Figure 3. Literal C_1 and C_2 are true in $\mathcal{I}$, hence there is only one rule node for rules (2) and (3). Neither $root(x)$ nor $\neg root(x)$ can be justified for all $x \in D$, hence they have been omitted.

There are four solutions that are subset-maximal with respect to rule nodes, namely the following rule node selections:

$$\{\langle R, \mathbf{u}\rangle, \langle root, \mathbf{u}\rangle, \langle C_2, \mathbf{t}\rangle\} \quad \text{(a)}$$
$$\{\langle R, \mathbf{f}\rangle, \langle C_2, \mathbf{t}\rangle\} \quad \text{(b)}$$
$$\{\langle R, \mathbf{t}\rangle, \langle C_2, \mathbf{t}\rangle\} \quad \text{(c)}$$
$$\{\langle C_1, \mathbf{t}\rangle, \langle C_2, \mathbf{t}\rangle\} \quad \text{(d)}$$

For each of these, multiple justification selections are possible (also shown in Figure 3). For C_1, we have to select justification (iv), but for C_2 we can choose from (v) or (vi) (but not both).

The objective is then not to maximize the number of selected rule nodes, but to minimize the expected grounding size. To obtain an estimate of the expected grounding size, the following conditions were taken into account:

- It should depend on the size of the grounding of the rule.
- Assigning multiple justifications to a rule should result in a lower estimate as the rule will only be grounded when all are false.

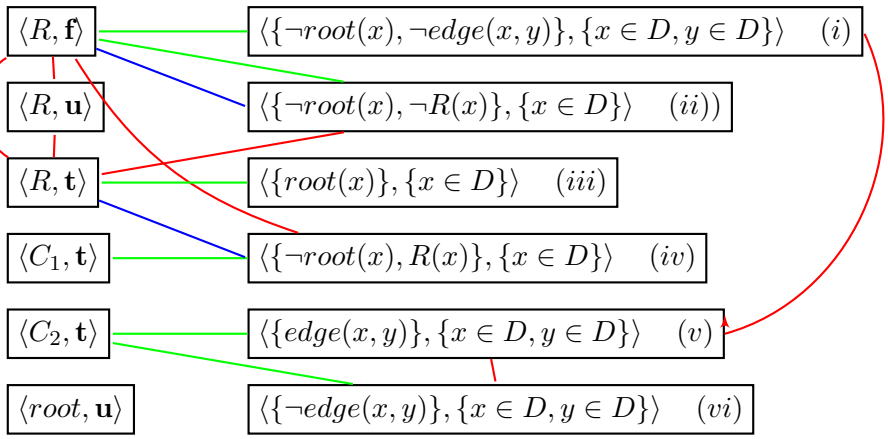

Figure 3: The graph that is part of the input to the optimal justification problem of Example 4.9. Rule nodes are shown on the left, justification nodes on the right; valid edges are shown in green, conflict edges in red and depends-on edges in blue. For readability, conflicts between justifications and unknown rule nodes are not shown.

- Shared variables result in less matching instantiations.
- In most practical applications, the number of false atoms far exceeds the number of true ones in any model. Hence, positive literals in a justification should have a higher cost than negative ones.

We approximate the expected grounding size with the function exp_{size} which takes as input a rule r, the selected type of justification (rule not selected (**n**), no justification (**u**), making the rule true (**t**) or false (**f**)) and the set of justifications J. It is then defined as

$$
\begin{aligned}
exp_{size}(r, \mathbf{n}, \emptyset) &= size(r) \\
exp_{size}(r, \mathbf{u}, \emptyset) &= size(r) \times 0.1 \\
exp_{size}(r, \mathbf{f}, J) &= size(r) \times \prod_{J_i \in J} (0.3 \times |\text{pos. lits.} \in J_i| \\
&\qquad + 0.1 \times |\text{neg. lits.} \in J_i|) \\
exp_{size}(r, \mathbf{t}, J) &= size(r) \times 0.3 \times \prod_{J_i \in J} (0.3 \times |\text{pos. lits.} \in J_i| \\
&\qquad + 0.1 \times |\text{neg. lits.} \in J_i|)
\end{aligned}
$$

Below, we define the function $size$, an approximation of final grounding size of a rule, which takes the local approach into account. For an existential

quantification and disjunction, we take the logarithm to reflect that it generally results in less grounding than universal quantification or conjunction in the same context, as Tseitin introduction can be used to ground it only partially.

$$\begin{aligned} size(L) &= 1 \\ size(\exists x \in D : \varphi) &= log(D) \times size(\varphi) \\ size(\forall x \in D : \varphi) &= D \times size(\varphi) \\ size(\phi_1 \wedge \ldots \wedge \varphi_n) &= \sum_{i \in [1,n]} size(\varphi_i) \\ size(\phi_1 \vee \ldots \vee \varphi_n) &= log(n) \times \frac{\sum_{i \in [1,n]} size(\varphi_i)}{n} \\ size(L \leftarrow \varphi) &= size(\varphi) + 1 \end{aligned}$$

Solutions to the optimal justification problem should minimize the term $\sum_{r \in \Delta_{\mathrm{d}}} exp_{size}(r, t(r), J(r))$ with $t(r)$ the type of justification and $J(r)$ the constructions assigned to r.

Example 4.10 (Continued from Example 4.9)**.** The size of rule C_1 is $1 + log(D) \times 2$, that of C_2 is $1 + D^2 \times log(2)$ and that of R is $D \times (1 + log(2) \times (1 + log(D) \times 2)/2$. Consider assigning justification (iv) to C_1: this results in an expected cost for that rule of $(1 + log(D) \times 2) \times 0.3 \times 1$ (as the construction relies on making R true). Additionally, it would force the grounding of the rule defining R, increasing the cost with the size of the rule for R. The optimal solution for the problem in Figure 3 is then rule node selection (a) with justification (vi) for C_2. Its cost is the sum of $(1 + D^2 \times log(2) \times 0.3$ (for justification (vi)) and $1 + log(D) \times 2$ (the expected size of the rule for C_1). With this solution, only the rule for C_1 has to be passed to the local approach.

To solve the optimal justification problem, IDP's optimization inference itself is applied to a (meta-level) declarative specification of the task.[3] For larger theories $\mathcal{T}$, the problem turns out to be quite hard, so two approximations were considered to reduce the search space. First, the number of selected justifications for any rule was limited to 2. Second, as the values of $size(r)$ can grow quite large, the approximation $\lceil \log(size(r)) \rceil$ is used (a standard approach). The resulting specification could be solved to optimality within seconds for all tested theories. During lazy model expansion, the global approach is applied in the initial phase when all Tseitin literals representing sentences in the original theory have been propagated true.

[3]The specification is part of IDP's public distribution.

5 Optimizations and Discussion

In this section, a number of optimizations is discussed that improve specific parts of the presented algorithms. First, we discuss heuristic choices affecting in the different algorithms. +Second, it is briefly discussed how the create an even smaller grounding and how to avoid confusing the search algorithm with too many spurious Tseitin variables. Finally, we discuss how to extend the approach to the logic FO$(\cdot)^{\texttt{IDP}}$, an extension of FO(ID), and to related inference tasks.

5.1 Heuristics

The algorithms leave room for a number of heuristic choices, that can have an important effect on the performance. We now briefly discuss these choices. As a guideline for our decisions, the following principles are used:

- Avoid leaving the search process without enough information to make an informed decision; for example avoid losing too much (unit) propagation or introducing too many Tseitin symbols.
- Prevent creating a grounding that is too large; this may for example happen as the result of a very long propagate-ground sequence.

Recall, the goal is not to create a minimal grounding, but to solve model expansion problems while avoiding a too large grounding.

In split_and_ground, when handling a disjunction or existential quantification, there is a choice on how many disjuncts to expand. If we expand one instantiation at a time for a rule $h \leftarrow \exists x \in D : P(x)$, as done in Algorithm 7 (lines 3 and 7), iterative application results in a ground theory

$$\begin{aligned} h &\leftarrow P(d_1) \vee T_1 \\ T_1 &\leftarrow P(d_2) \vee T_2 \\ T_2 &\leftarrow P(d_3) \vee T_3 \\ &\vdots \\ T_n &\leftarrow \exists x \in D \setminus \{d_1, d_2, \ldots, d_n\} : P(x). \end{aligned}$$

If we adapt the algorithm to introduce n elements at a time, we obtain

$$\begin{aligned} h &\leftarrow P(d_1) \vee \ldots \vee P(d_n) \vee T \\ T &\leftarrow \exists x \in D \setminus \{d_1, d_2, \ldots, d_n\} : P(x). \end{aligned}$$

In the experiments, adding 10 instantiations of an existential quantification and 3 disjuncts of a disjunction at a time turned out to give the best balance between introducing too many Tseitin atoms and grounding too much.

The thrashing behaviour that leads to the above-mentioned propagate-grounding chains for a disjunction or existential quantification can have several origins. One if them is the truth value assigned to first-time choice atoms by the search algorithm. For example, in the SAT-solver Minisat (used in the IDP system), initially each atom is assigned **f**. This causes the above sentence $\exists x \in D : P(x)$ to be fully grounded (as a conflict only arises when the full disjunction is grounded). To remedy this, two search-related heuristics are changed. First, the initial truth value is randomized (still favoring false), by assigning true initially with probability 0.2. Second, search algorithms typically *restart* after an (ever-increasing) threshold on the number of conflicts, sometimes caching the truth value assigned to atoms (*polarity caching*). This allows to take learned information into account in the search heuristic while staying approximately in the same part of the search space. In case of lazy grounding, we might want to jump to another part of the search space when we come across long propagate-ground sequences. To this end, we introduce the concept of *randomized restarts*, which take place after an (ever-increasing) threshold on the number of times Δ_g is extended and randomly flipping some of the cached truth values. The initial threshold on the number of extensions is 100; it is doubled after each restart.

In addition, build_djust always returns *false* if it is estimated that the formula has a small grounding. Indeed, grounding them can help the search. In our experiments, a formula is considered small if its (estimated) grounding size is below 10^4 atoms. The same strategy is used in split_and_ground: whenever the formula to which one_step_ground would be applied is sufficiently small, ground is applied instead, to completely ground the remaining formula.

5.2 Other issues

We briefly discuss a number of further considerations that have an important effect on the behaviour of the algorithm, but for which a deeper discussion is out of the scope of this text.

Propagation guarantees. In build_djust, we can vary how many domain element/disjuncts to select for justifying disjunctive knowledge (currently the whole domain, respectively one disjunct). Selecting only one object will result in minimal grounding, but also only guarantees *consistency*: if

the justification becomes false, the associated sentence might already be false. Consequently, if its grounding had already been part of Δ_g, it might have resulted in earlier propagation. There is an important subclass of justifications that guarantee that propagation is not delayed.

Definition 5.1 (unit-propagation-safe)**.** A justification J of a rule r of the form $h \leftarrow \varphi$ is *unit-propagation-safe* if (**i**) the head is true (false) in $\mathcal{I}$ and the justification for h ($\neg h$) can be partitioned in two valid justification or (**ii**) the head is unknown in $\mathcal{I}$ and J is to be a valid justification for φ and contains $\neg h$ or is a valid justification for $\neg\varphi$ and contains h.

Informally, propagation might get delayed if a justification is still available, but it is the only way in which the associated sentence can still become true. By guaranteeing two disjoint justifications are available, we are sure no unit propagation (or detection of a conflict) might be delayed. This is in fact a generalization of the two-watched literal scheme (Moskewicz, Madigan, Zhao, Zhang, & Malik, 2001), an important optimization in SAT-solvers: for each clause, it is sufficient to maintain 2 non-false literals and only check propagation for the clause when either one becomes false. To incorporate propagation-safe justifications in our algorithm, some changes have to be made, such as adapting build_djust to only return such justifications. Also, we cannot just delay grounding of a rule as long as its head is not relevant (as its body might already be true or false). It is part of future work to experiment with this type of justifications. An obvious disadvantage is that it results in earlier grounding, but this might be offset by improved search. On the other hand, note that the presented algorithms lose propagation only **once**: when the justification becomes false, the relevant sentence is grounded, hence afterwards it is checked for propagation as it has become part of Δ_g.

Cheap propagation checks. In lazy_mx, for each assigned literal it is checked whether it is defined in Δ_d and whether it violates any justifications. To implement this cheaply, a mapping is maintained from literals in Δ_g to whether they are defined in Δ_d and in which justifications its negation occurs. This mapping is created whenever a literal is first added to Δ_g and maintained whenever justifications change. Consequently, the performance of the search loop is unaffected as long as literals are assigned for which the mapping is empty.

Stopping early. In Algorithm 2, we took the standard stopping criterion used in most search algorithms (Line 14): to stop when $\mathcal{I}$ is two-valued (on

$T_g \cup \Delta_g$) and no conflict has been found, as it is now certainly a model of the theory. In our setting, we would prefer to stop earlier, ideally as soon as $\mathcal{I}$ is a total justification for $P_\mathcal{T}$, which is correct as shown in Corollary 3.5. Otherwise, we might ground more than necessary. Indeed, for example assigning literals defined in Δ_d by rules without a justification results in violated rules and possibly further grounding, even when the literal was not necessary for the justification of $P_\mathcal{T}$. Our algorithms only chooses literals that are watched by some formula/rule. It guarantees that if all watched literals have been assigned and no conflict ensues, that the current partial structure satisfies the ground theory and hence search can stop.

Alternative interleaving of grounding and search. Grounding is applied during the search process as soon as unit propagation has taken place. The result is a focus on the current location in the search space, but with the danger of grounding too much if there is no solution in that part of the space. Alternatively, we could apply the opposite strategy, namely to ground as late as possible: only apply additional grounding when the search algorithm terminates without ever having found a model in an acceptable default state. Such a strategy is well-known from the fields of incremental proving and planning, where the domain (number of time steps) is only increased after search over the previous, smaller bound has finished. The latter strategy guarantees a minimal grounding. A prototype of this strategy has been implemented in IDP with good results on planning problems.

Approximate justifications. In some cases, build_djust cannot find a valid justification for a large formula because a few literals are already false in $\mathcal{I}$. For example for a formula $\forall x \in D : P(x)$, build_djust returns *false* if at least one P is false. Instead, we adapt build_djust with a heuristic check on the number of expected violations. If it is small enough, the justification is still returned. Naturally, we are then required to check whether there are any real violations, by querying the justification formula over $\mathcal{I}$, and apply lazy_ground to them.

5.3 Extension to FO$(\cdot)^{\text{IDP}}$

So far, we have described a lazy model expansion algorithm for function-free FO(ID). However, FO$(\cdot)^{\text{IDP}}$, the knowledge-base language of the IDP system, supports a richer input language that includes partial functions, aggregates, arithmetic and types. More work is needed to fully exploit the potential of lazy grounding for these extensions. However, a straightforward

implementation is to slightly adapt build_djust (Algorithm 5): the literal case is extended to return *false* when the literal is not a function-free literal and an extra case, also returning *false*, is added to handle expressions that do not belong to function-free FO(ID). For example, given a rule $h \leftarrow \forall x : P(x) \vee Q(f(x))$, $P(x)$ can be used in a justification but $Q(f(x))$ cannot. This is the approach used in the experiments of Section 6.

5.4 Related inference tasks

The bulk of the paper focuses on model expansion (MX) for FO(ID) theories $\mathcal{T}$, for which solutions are structures which are two-valued on $voc(\mathcal{T})$. Often, one is only interested in a small subset of the symbols in $voc(\mathcal{T})$. This is for example the case for model generation for $\exists$SO(ID), the language which extends FO(ID) with existential quantification over relations. An $\exists$SO(ID) problem $\exists P_1, \ldots, P_n : \mathcal{T}$ with an initial structure $\mathcal{I}$, relation symbols $P_1, \ldots, P_n$, and $\mathcal{T}$ an FO(ID) theory, can be solved by model generation for the FO(ID) theory $\mathcal{T}$ with initial structure $\mathcal{I}$ and by dropping the interpretation of the symbols $P_1, \ldots, P_n$ in the models. Another example is query evaluation for FO(ID): given a theory $\mathcal{T}$, an initial structure $\mathcal{I}$ and a formula φ with free variables $\overline{x}$ (all in FO(ID)), the purpose of evaluating the query $\langle \mathcal{T}, \mathcal{I}, \varphi \rangle$ is to find assignments of domain elements $\overline{d}$ to $\overline{x}$ such that a model of $\mathcal{T}$ exists that expands $\mathcal{I}$ and in which $\varphi[\overline{x}/\overline{d}]$ is true. To solve it by model expansion in FO(ID), a new predicate symbol T is introduced and answers to the query are tuples of domain elements $\overline{d}$ such that $T(\overline{d})$ is true in a model of the theory $\mathcal{T}$ extended with the sentence $\exists \overline{x} \in \overline{D} : T(\overline{x})$ and the definition $\{\forall \overline{x} \in \overline{D} : T(\overline{x}) \leftarrow \varphi\}$.

In both cases, approaches using (standard) model expansion compute a total interpretation and afterwards drop all unnecessary information, which is quite inefficient. Lazy model expansion can save a lot of work by only partially grounding the theory. However, once a model is found for the grounded part, the justifications and the remaining definitions are used to expand the structure to a model of the full theory. Although this expansion is obtained in polynomial time, it is still inefficient when afterwards a large part of the model is dropped.

To remedy this, we define a variant of the model expansion task, denoted *restricted* MX. Restricted MX takes as input a theory $\mathcal{T}$, a structure $\mathcal{I}$ and an additional list of symbols O, called *output symbols*.[4] Solutions are then structures M which are two-valued on all symbols in O and for which

[4] Within the ASP community, they are sometimes referred to as "show" predicates.

an expansion exists that extends $\mathcal{I}$ and is a model of $\mathcal{T}$. Adapting lazy grounding to solve restricted MX can be done through an analysis of which justifications need not be added (completely) to the structure, splitting Δ_{gd} into multiple definitions and only evaluating those defining output symbols or symbols on which those depend (using a stratification argument).

The above-mentioned inference tasks can be cast trivially to restricted MX problems and lazy restricted MX then greatly improves the efficiency with respect to ground-and-solve, as shown in the experimental section.

The extension of FO(ID) with *procedurally interpreted* symbols (De Cat et al., 2014) provides another class of interesting problems. Such predicate symbols have a fixed interpretation, but to know whether a tuple belongs to the predicate, a procedural function has to be executed. Such an approach provides a clean way to combine declarative and procedural specifications. Consider for example a symbol $isPrime(\mathbb{N})$ that is interpreted by a procedure which executes an efficient prime-verification algorithm and returns true iff the given argument is prime. We are generally not interested in the complete interpretation of $isPrime$, so it can be cast as a restricted MX problem with $isPrime$ not in O. Solving such a problem using lazy grounding then has the benefit of only executing the associated function *during* search for relevant atoms $isPrime(d)$. Also for this task, we show an experimental evaluation in the next section.

6 Experiments

The IDP system has a state-of-the-art model expansion engine, as can be observed from previous Answer-Set Programming competitions ((Denecker, Vennekens, Bond, Gebser, & Truszczyński, 2009), (Calimeri, Ianni, & Ricca, 2012) and (Alviano et al., 2013)). The lazy model expansion algorithms presented in this paper were implemented in the IDP system, by extending the existing algorithms (De Cat, Bogaerts, Denecker, & Devriendt, 2013a). The implementation is preliminary in the sense that direct justifications are only derived for sentences, but not for definitions. However, the grounding of rule instances can be delayed until their head is assigned and Tseitin introduction can be applied to disjunctive bodies. As non-inductive definitions can be equivalently expressed as their completion, this does not significantly restrict the applicability of the implementation.

We tested three different setups: IDP with the standard ground-and-solve approach (referred to as `g&s`), IDP with lazy model expansion (`lazy`) and the award-winning ASP system Gringo-Clasp (`ASP`). All experiments

were run on an 64-bit Ubuntu 13.10 system with a quad-core 2.53 Ghz processor and 8 Gb of RAM. A timeout of 1000 seconds and a memory limit of 3 Gb was used; out-of-time is indicated by T, out-of-memory by M. We used IDP version 3.2.1-lazy, Gringo 3.0.5 and Clasp 2.1.2-st.[5]

The section is organized as follows. In Section 6.1, we evaluate the overhead of completely grounding a theory using the presented approach. In Section 6.2, we evaluate the effect of lazy grounding on a number of benchmarks of the ASP competition. Afterwards, in Section 6.3, a number of additional properties of the presented algorithms are demonstrated.

6.1 Lazy grounding complexity

To quantify the overhead of lazy model expansion over the standard grounding algorithm, we evaluated the time necessary to completely ground a theory over a given structure in both IDP setups. To this end, we minimally adapted the lazy grounding implementation to immediately add literals in the ground theory or in a justification to the q_{ch} queue and by setting the threshold under which a formula is considered "small" to 0 atoms. As discussed previously, we expected the (naive) incremental querying of justifications to be a bottleneck.

We devised six benchmarks to test various aspects of the novel algorithm that set it apart from non-incremental grounding. The tested aspects are:

1. Existential quantification using Tseitin transformation.
2. Definitions without justification for the body.
3. Sentences without disjunctive justifications.
4. The effect of intelligent grounding techniques that derive smaller bounds on variable instantiations (Wittocx, Mariën, & Denecker, 2010b; Wittocx, Denecker, & Bruynooghe, 2013b). We distinguish between standard "grounding without LUP" (for Lifted Unit Propagation), which derives bounds for a quantified variable based only on its subformula, and "grounding with LUP", which also incorporates information from the rest of the theory.
5. Nested universal quantification for a sentence of the form $\forall \overline{x} : P(\overline{x}) \Rightarrow \forall \overline{y} : Q(\overline{x}, \overline{y})$. Recall, the implementation of the local approach only derives justifications on one quantification level. As a result, first the

[5] Benchmarks, experimental data and complete results are available at dtai.cs.kuleuven.be/krr/files/experiments/lazygrounding2014experiments.tar.gz.

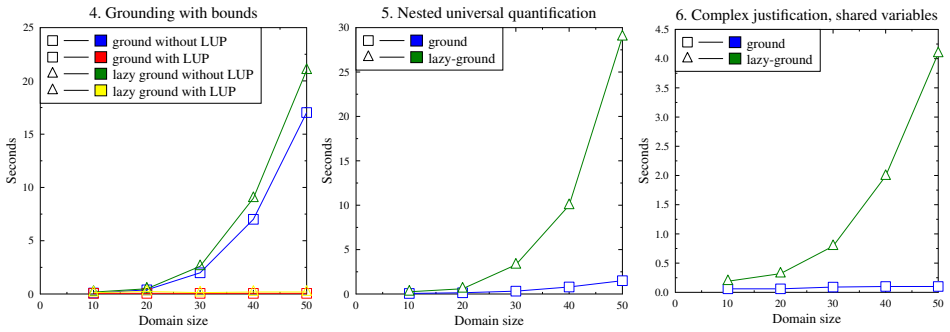

Figure 4: Time overhead of lazy grounding over ground-and-solve when completely grounding the input theory, for benchmarks 4, 5 and 6.

justification is found which sets P false; when this wakes up, it results in splitting the formula and a new justification that sets Q true.

6. Sentences with disjunctive justifications with shared variables.

Experiments were done with domains of size 10, 20, 30, 40 and 50.

In all experiments, the overhead for the time required to solve the initial optimization problem (for the global approach) was always around 0.02 second, so in itself negligible. The results for the first three experiments are not shown as the differences were negligible, as expected. The results for the latter three are shown in Figure 4. As expected, there is a significant overhead to handle more complex justification formulas, especially when variables are shared (in which case significant query optimization is possible). However, these results are a worst-case view of the overhead: in practice, we expect that often, constructing the full grounding will not be necessary in the first place. Still, it is an important part of future research to improve the querying algorithm to reduce it to an acceptable overhead.

6.2 ASP competition benchmarks

Second, we selected benchmarks from previous ASP competitions to evaluate the lazy grounding algorithm in a more realistic setting. For most of the available benchmarks, the encoding can usually be sufficiently optimized so that the grounding is not too large. Hence, to show the improvements of lazy grounding over ground-and-solve, we focused on benchmarks in which a natural, predicate encoding can easily result in a large grounding. We also selected some benchmarks that are known to be hard, to evaluate the effect

of lazy model expansion on search. We selected the following problems (see the competition websites for complete descriptions):

- `Reachability`: Given a directed graph, determine whether a path exists between two given nodes.
- `Labyrinth`: A planning problem where the aim is to manipulate a graph such that the goal node becomes reachable from the start node.
- `Packing`: Given a 2-D rectangle and a number of squares, fit all squares into the grid without overlaps.
- `Disjunctive Scheduling`: Schedule a number of actions with a given earliest start and latest end time with additional constraints on precedence and disjointness.
- `Sokoban`: A planning problem where a robot has to push a number of blocks to goal positions, constrained by a 2-D maze.
- `Graph Colouring`: Given a graph, assign a colour to each node (from a given set of colours), such that no directly connected nodes have the same colour.
- `Stable Marriage`: Given a set of men and women and a set of match preferences, find an assignment that is "stable": no swap results in a better or equal match.
- `Hydraulic Planning`: A planning problem where the task is to turn on a number of valves in a pipe system in the proper order such that some goal nodes become pressurized.

For each of these, we randomly selected 10 of the 2011 competition instances (2013 for `Stable Marriage`). For `Stable Marriage`, `Graph Colouring` and `Reachability`, we based our encodings on the available ASP-Core-2 encodings. For `Packing`, `Disjunctive Scheduling`, `Hydraulic Planning` we constructed a natural IDP encoding and made a faithful translation to ASP. For the more complex benchmarks of `Labyrinth` and `Sokoban`, we used the original IDP and Gringo-Clasp specifications sent in to the 2011 competition. For the lazy model expansion, we replaced cardinality expressions by their FO encoding as for the former no justifications are derived yet; this also increases the size of the full grounding. Note that the encoding used is not necessarily the most efficient one. Specifically for `Disjunctive Scheduling` and `Hydraulic Planning`, more efficient encodings are available, using uninterpreted functions and an encoded algorithm respectively.

	# solved			avg. time (sec.)		
benchmark	g&s	lazy	ASP	g&sIDP	lazy IDP	ASP
Sokoban	9	5	**10**	370	1538	**22**
Disjunctive Sched.	4	**10**	2	1933	**92**	2400
Packing	9	**10**	1	352	**51**	2704
Labyrinth	**6**	5	**6**	1286	1851	**1015**
Reachability	1	**10**	4	2706	**7**	1807
Stable Marriage	0	**10**	5	3000	**350**	1563
Graph Colouring	**7**	3	4	**1013**	2103	1842

Table 1: The number of solved instances for the ASP benchmarks (10 instances each) and the average time taken.

For Hydraulic Planning, none of the instances could be solved by any setup. For the other benchmarks, the number of solved instances and average time are shown in Table 1; the average grounding size for the IDP setup is shown in Table 2. For unsolved instances, we took three times the time threshold as time and 10^{10} as grounding size in case of memory overflow.[6] Memory overflows only happened in Reachability (7 times for g&s), Labyrinth (4 times for g&s, once for lazy) and Packing (once for g&s); all other unsolved instances were caused by a time-out.

The results show that lazy model expansion solved more instances than the other setups in four out of seven cases. In those cases, the problems also got solved significantly below the time threshold. In five out of seven cases, the (final) grounding size was smaller for lazy model expansion, even orders of magnitude in three cases. For Sokoban, Labyrinth and Graph Colouring, lazy model expansion was outperformed by ground-and-solve, indicating that the loss of information outweighed the gain of grounding less up-front. E.g., for Sokoban, the lazy final grounding size was even higher than for g&s(possible due to the FO encoding of cardinalities), indicating that a large part of the search space was explored. For Stable Marriage, the relatively small difference in grounding size leads us to believe that the different search heuristic was the main factor and not lazy grounding itself.

For hydraulic planning (for which no instance was solved), we observed that for ground-and-solve, the computation went out of memory, while lazy model expansion was always able to start searching, even though it was not able to find any solutions. A similar observation was made during ex-

[6]IDP has automatic symmetry breaking, the cause of the difference between g&s and ASP for Graph Colouring.

benchmark	grounding size (# atoms) g&s	lazy
Sokoban	$\mathbf{2.5}^4$	7.8^5
Disjunctive Sched.	6.8^7	$\mathbf{3.5}^6$
Packing	1.0^9	$\mathbf{1.4}^7$
Labyrinth	4.0^9	$\mathbf{2.0}^9$
Reachability	7.1^9	$\mathbf{1.5}^4$
Stable Marriage	4.4^7	$\mathbf{1.5}^7$
Graph Colouring	1.1^4	1.6^4

Table 2: The average grounding size for the ASP benchmarks for both IDP setups. For the lazy setup, the size of the final ground theory was taken.

periments with the Airport Pickup ASP-2011 benchmark, a fairly standard scheduling problem (transporting passengers by taxis taking into account fuel consumption) except that no upper bound on time is provided. Hence any ground-and-solve approach would need to construct an infinite grounding. Applying straightforward lazy model expansion also resulted in a grounding that was too large. However, using a prototype in which additional grounding is done when unsatisfiable has been found (see Section 5.2), IDP solved one out of ten instances. For the others, grounding was not the problem, but the search heuristic took too long at each of the time intervals $1..n$ considered to get up to a sufficient n to solve the problem (even with the standard search heuristic).

The presented results show that, although often beneficial, lazy model expansion can be a considerable overhead for some hard search problems. Therefore, the logical step is with the current algorithms is then a *portfolio approach*, which runs both the ground-and-solve and lazy model expansion in tandem. We expect such an approach to come close to the best solution in all presented benchmarks. However, on all the problems considered, lazy model expansion could start search much earlier than ground-and-solve, even though it got lost more often during search. This leads us to believe that combining lazy model expansion with a better heuristic (or possible a user-specified one) will allow it to demonstrate full capabilities.

6.3 Specific experiments

Next to the ASP competition benchmarks, some experiments were conducted using constructed benchmarks/instances to illustrate specific properties of the lazy grounding algorithm.

The first part of Table 3 shows the results of scalability experiments. For each of the benchmarks `Packing`, `Sokoban` and `Disjunctive Scheduling`, we selected a very simple instance and evaluated the effect of increasing the domain size (either time or grid size) by orders of magnitude. The results shows that for each of the instances, lazy model expansion scales much better than ground-and-solve, both for IDP and Gringo-Clasp and for satisfiable and unsatisfiable instances. However, for `Disjunctive Scheduling` the solving time still increases significantly. The reason is that the lazy heuristics are still naive and make uninformed choices too often.

The second part of the table shows results for some crafted benchmarks:

- `Dynamic reachability`, the example described in 3.3.
- Lazy evaluation of `procedurally interpreted` symbols, using a simple theory over the prime numbers. As described in Section 5.4, a predicate symbol $isPrime/1$ is interpreted by a procedure that returns true if the argument is prime.
- A predicate encoding of a `function` with a huge domain (note that a state-of-the-art encoding using aggregates or uninterpreted functions instead of a predicate encoding would perform better).
- A `modelgeneration`-like application, where the aim is to also find the domain of the structure.

For each one, a faithful ASP encoding was constructed. The results show a significant improvement of lazy model expansion over ground-and-solve on all examples.

Closer to inherent complexity? During the modeling phase of an application, different encodings are typically tested out, in an attempt to improve performance or to locate bugs. While modeling our experimental benchmarks, we noticed that simplifying a theory by dropping constraints often resulted in a dramatic reduction in the time lazy model expansion took to find a model. Standard model expansion, on the other hand, was much less affected by such simplifications. In our opinion, this observation, while hardly definitive evidence, is another indication that the presented algorithms are able to derive justifications for parts of a theory that can be satisfied cheaply. In that way, the approach is able to distinguish better between problems which are inherently difficult and problems which would just have a large grounding.

benchmark	lazy	g&s	ASP
packing-10	0.2	2.0	0.1
packing-25	0.3	2.0	0.1
packing-50	1.1	10.03	5.8
sokoban-10^3	0.31	0.3	0.1
sokoban-10^4	0.5	20.0	1.1
sokoban-10^5	2.6	T	68.0
disj-sched-sat-10^3	0.39	0.49	0.07
disj-sched-sat-10^4	13.04	16.05	17.44
disj-sched-sat-10^5	164.18	M	M
disj-sched-unsat-10^3	0.24	0 49	0.09
disj-sched-unsat-10^4	4.11	16.04	19.85
disj-sched-unsat-10^5	164.2	M	M
dynamic reachability	0.18	M	M
procedural	1.24	M	M
function	0.79	M	M
modelgeneration	0.19	M	M

Table 3: Additional experiments on crafted benchmarks, one instance each.

7 Related work

Lazy model expansion offers a solution for the blowup of the grounding that often occurs in the ground-and-solve model expansion methodology for FO(ID) theories. ASP and Sat-Modulo-Theories techniques also process theories that can have a large grounding; the constraint store of CP and Mixed Integer Programming and the clauses of SAT can be considered the equivalent of a grounded theory (they are often derived from quantified descriptions such as "$c_i < c_j$ **for all** i and j for which ...") and can also become very large. Various authors have reported a blowup problem in these related paradigms (Lefèvre & Nicolas, 2009; Ge & de Moura, 2009) and a multitude of techniques has been developed to address it. We distinguish four general approaches.

First, concerning grounding up-front, research has been done towards *reducing the size of the grounding* itself through (**i**) *static analysis* of the input to derive bounds on variable instantiations (Wittocx et al., 2010a, 2013a), (**ii**) techniques to *compile* specific types of sentences into more compact ground sentences (Tamura et al., 2009; Metodi & Codish, 2012), (**iii**) detect parts that can be evaluated polynomially (Leone et al., 2006; Gebser

et al., 2011; Jansen et al., 2013) and (**iv**) detect parts that are not relevant to the task at hand (e.g., in the context of query problems) (Leone et al., 2006). Naturally, each of these approaches can be used in conjunction with lazy grounding to further reduce the size of the grounding. In IDP for example, lazy grounding is already combined with (**i**) and (**iii**).

Second, the size of the grounding can be reduced by *enriching* the language the search algorithm supports. For example, ASP solvers typically support ground aggregates (interpreted second-order functions such as cardinality or sum that take sets as arguments), and CP and SMT solvers support (uninterpreted) functions. More recently, the Constraint-ASP paradigm was developed (Ostrowski & Schaub, 2012), that integrates ASP and CP by extending the ASP language with *constraint* atoms. These are interpreted as constraints in a CSP problem and can thus be handled using CP techniques. Various CASP solvers are already available, such as Clingcon(Ostrowski & Schaub, 2012), Ezcsp(Balduccini, 2011), Mingo (Liu, Janhunen, & Niemelä, 2012) and Inca (Drescher & Walsh, 2011) and is also integrated into IDP (De Cat et al., 2013b). Inca and IDP in fact implement Lazy Clause Generation (Stuckey, 2010), an optimized form of lazy grounding for specific types of constraints. The language HEX-ASP (Eiter et al., 2005) also extends ASP, this time with *external* atoms that represent (higher-order) external function calls.

Third, *incremental grounding* is a technique well-known from model generation, theorem proving and planning. For these tasks, the domain is typically not fixed in advance, but part of the structure being sought, such as the number of time steps in a planning problem. Incremental grounding then works by grounding the problem for an initial guess of (the number of elements in) the domain. Afterwards, search is applied; if no model was found, the domain is extended and more grounding is done. This is iterated until a model is found or a bound on the maximum domain size is hit (if one is known). This technique is applied, e.g., in the prover Paradox (Claessen & Sörensson, 2003) and the ASP solver IClingo (Gebser et al., 2008).

Fourth, and closest to lazy grounding itself, is a large body of research devoted to delaying the grounding of specific types of expressions until necessary (for example until they result in propagation). Propagation techniques on the first-order level that delay grounding until propagation ensues have been researched within ASP (Lefèvre & Nicolas, 2009; Palù et al., 2009; Dao-Tran et al., 2012) and within CP (Stuckey, 2010). Such techniques can be used in conjunction with lazy grounding as they derive more intelligent justifications for specific types of constraints than presented here. For example, in (Dao-Tran et al., 2012) the authors present an efficient al-

gorithm for bottom-up propagation in a definition. Within SMT, various theory propagators work by lazily transforming their theory into SAT, such as (Bruttomesso, Cimatti, Franz?n, Griggio, Hanna, Nadel, Palti, & Sebastiani, 2007) for the theory of Bit Vectors. Ge et al. (Ge & de Moura, 2009) investigated quantifier handling by combining heuristic instantiation methods with research into decidable fragments of FO theories, as these can be efficiently checked for models. In (Saptawijaya & Pereira, 2013), an abduction framework is extended to lazily generate part of the relevant sentences. In search algorithms themselves, justifications (or *watches*) are used to derive when a constraint will not result in propagation or is already satisfied, and hence need not be checked in the propagation phase. In (Nightingale et al., 2013), it is shown how maintaining (short) justifications can significantly reduce the cost of the propagation phase.

In fact, a well-known technique already exists that combines search with lazy instantiation of quantifiers, namely *skolemization*, where existentially quantified variables are replaced by newly introduced function symbols. Universal quantifications are handled by instantiating them for those introduced function symbols. Reasoning on consistency can, e.g., be achieved by congruence closure algorithms, capable of deriving consistency without effectively assigning an interpretation to the function symbols. These techniques are used in Tableau theorem proving (Hähnle, 2001) and SMT solvers (Detlefs, Nelson, & Saxe, 2005). Formula (Jackson, Bjorner, & Schulte, 2013) interleaves creating a ground program and giving it to an SMT solver, iterating when symbolic guesses proved to be wrong. Skolemization-based techniques typically work well in case only a small number of constants needs to be introduced, but have difficulty in case the relevant domain is large. One can also see that lazy grounding (with support for function symbols) could incorporate skolemization by adapting the rules for grounding existential and universal quantification. We expect skolemization to be complementary to lazy grounding, but an in-depth investigation is part of future work.

In the field of probabilistic inference, several related techniques have been developed that also rely on lazy instantiation. First, the Problog system uses a form of static dependency analysis to ground a (probabilistic) program in the context of a given query, by constructing all possible ways to derive the query in a top-down fashion (Kimmig et al., 2011). Second, so-called *lazy inference*, applied e.g. in *LazySAT* (Singla & Domingos, 2006), exploits the fact that, for the considered inference, a (fixed) *default* assumption exists under which an expression certainly does not contribute to the probabilities. Hence, expressions for which the assumption certainly holds do not have to be considered during search. Third, *cutting place inference* (Riedel, 2009)

applies lazy inference in an interleaved setting, only constructing the part of the program for which the assumptions are not satisfied.

8 Future work

In future, several aspects of the presented work will be investigated further.

One aspect is extending support to lazy ground more complex expressions, including aggregate expressions and (nested) function terms. Consider for example the sentence $(\sum_{x \in D \text{ and } P(x)} f(x)) > 3$, which expresses that the sum of terms $f(d)$ for which the atom $P(d)$ is true, with $d \in D$, P a predicate and f a function, should be larger than 3. One can observe that it is not necessary to ground the whole sentence up-front. E.g., if f maps to the natural numbers (hence positive), the set $\{P(d_1), f(d_1) > 3\}$ is a minimal justification. Even if no easy justification can be found, we can suffice by grounding only part of the sentence and delay the remainder. E.g., we can create the ground sentence $(\sum_{P(d_1)} f(d_1)) > 3 \vee T$, with T a Tseitin symbol defined as $(\sum_{P(d_1)} f(d_1)) + (\sum_{x \in D \setminus d_1 \text{ and } P(x)} f(x)) > 3$. Indeed, in any model of the sentence in which T is false, the original inequality is satisfied.

A second aspect is whether there are advantages to grounding earlier, for example to guarantee no propagation is lost, or grounding later, possibly reducing the size of the grounding even more. For example, consider the sentences $P \Rightarrow \phi$ and $\neg P \Rightarrow \psi$, with ϕ and ψ both large formulas for which no justification was found. Instead of grounding at least one of the sentences, we might add P to the list of atoms the search algorithm has to assign and only ground either of the sentences when P has been assigned a value (it might even be that unsatisfiable is detected before grounding either one).

Lastly, what about lazy *forgetting*? As the ground theory is extended when making the structure more precise, the ground theory could be reduced again during backtracking. By storing the justification violations that caused which grounding, we can derive which grounding can be forgotten again if the violation is no longer problematic (e.g., after backtracking). For this, an algorithm needs to be developed which tracks grounding/splitting dependencies between rules given their justifications. This closely resembles techniques used in tableau theorem proving and SMT, where the theory at hand can be compacted when moving to a different part of the search space.

9 Conclusion

Solvers used in the domains of SAT, SMT and ASP are often confronted with problems that are too large to ground. Lazy model expansion, the technique described in this paper, interleaves grounding and search in order to avoid the grounding bottleneck. The technique builds upon the concept of a justification, a deterministic recipe to extend an interpretation such that it satisfies certain constraints. A theoretical framework has been developed for lazy model expansion for the language FO(ID) and algorithms have been presented to derive and maintain such justifications and to interleave grounding with state-of-the-art CDCL search algorithms. The framework aims at bounded model expansion, in which all domains are finite, but is also an initial step towards handling infinite domains efficiently. Experimental evaluation has been provided, using an implementation in the IDP system, in which lazy model expansion was compared with a state-of-the-art ground-and-solve approach. The experiments showed considerable improvement over ground-and-solve in existing benchmarks as well as in new applications. The main disadvantage is the less-informed search algorithm, caused by the delay in propagation and the introduction of (more) additional symbols. This leads us to believe a portfolio approach, heuristically selecting between ground-and-solve and lazy model expansion, would result in an overall improvement.

References

Alviano, M., Calimeri, F., Charwat, G., Dao-Tran, M., Dodaro, C., Ianni, G., Krennwallner, T., Kronegger, M., Oetsch, J., Pfandler, A., Pührer, J., Redl, C., Ricca, F., Schneider, P., Schwengerer, M., Spendier, L. K., Wallner, J. P., & Xiao, G. (2013). The fourth answer set programming competition: Preliminary report. In Cabalar, P., & Son, T. C. (Eds.), *LPNMR*, Vol. 8148 of *LNCS*, pp. 42–53. Springer.

Apt, K. R. (2003). *Principles of Constraint Programming.* Cambridge University Press.

Balduccini, M. (2011). Industrial-size scheduling with ASP+CP. In *LPNMR*, pp. 284–296.

Baral, C. (2003). *Knowledge Representation, Reasoning, and Declarative Problem Solving.* Cambridge University Press, New York, NY, USA.

Bruttomesso, R., Cimatti, R., Franz?n, A., Griggio, A., Hanna, Z., Nadel, E., Palti, A., & Sebastiani, R. (2007). A lazy and layered SMT(BV) solver for hard industrial verification problems. In *In Computer Aided Verification (CAV), LNCS*. Springer.

Calimeri, F., Ianni, G., & Ricca, F. (2012). The third open answer set programming competition. *CoRR*, *abs/1206.3111*.

Chen, W., & Warren, D. S. (1996). Tabled evaluation with delaying for general logic programs. *Journal of the ACM*, *43*(1), 20–74.

Claessen, K., & Sörensson, N. (2003). New techniques that improve MACE-style model finding. In *MODEL*.

Dao-Tran, M., Eiter, T., Fink, M., Weidinger, G., & Weinzierl, A. (2012). Omiga : An open minded grounding on-the-fly answer set solver. In del Cerro, L. F., Herzig, A., & Mengin, J. (Eds.), *JELIA*, Vol. 7519 of *LNCS*, pp. 480–483. Springer.

De Cat, B., Bogaerts, B., Bruynooghe, M., & Denecker, M. (2014). Predicate logic as a modelling language: The IDP system. *CoRR*, *abs/1401.6312*.

De Cat, B., Bogaerts, B., Denecker, M., & Devriendt, J. (2013a). Model expansion in the presence of function symbols using constraint programming. In *IEEE 25th International Conference on Tools with Artificial Intelligence, ICTAI 2013*, pp. 1068–1075.

De Cat, B., Bogaerts, B., Denecker, M., & Devriendt, J. (2013b). Model expansion in the presence of function symbols using constraint programming.. In *IEEE 25th International Conference on Tools with Artificial Intelligence, ICTAI 2013* (De Cat et al., 2013a), pp. 1068–1075.

De Cat, B., Denecker, M., & Stuckey, P. (2012). Lazy model expansion by incremental grounding. In Dovier, A., & Santos Costa, V. (Eds.), *Proceedings of the 28th International Conference on Logic Programming - Technical Communications (ICLP'12)*, Vol. 17 of *LIPIcs*, pp. 201–211. Schloss Daghstuhl - Leibniz-Zentrum fuer Informatik.

Delgrande, J. P., & Faber, W. (Eds.). (2011). *Logic Programming and Nonmonotonic Reasoning - 11th International Conference, LPNMR 2011, Vancouver, Canada, May 16-19, 2011. Proceedings*, Vol. 6645 of *LNCS*. Springer.

Denecker, M. (1998). The well-founded semantics is the principle of inductive definition. In Dix, J., del Cerro, L. F., & Furbach, U. (Eds.), *JELIA*, Vol. 1489 of *LNCS*, pp. 1–16. Springer.

Denecker, M., Bruynooghe, M., & Marek, V. W. (2001). Logic programming revisited: Logic programs as inductive definitions. *ACM Transactions on Computational Logic (TOCL)*, *2*(4), 623–654.

Denecker, M., & De Schreye, D. (1992). Justification semantics: A unifying framework for the semantics of logic programs. Tech. rep. 157, Department of Computer Science, K.U.Leuven.

Denecker, M., & De Schreye, D. (1993). Justification semantics: A unifying framework for the semantics of logic programs. In Pereira, L. M., & Nerode, A. (Eds.), *LPNMR*, pp. 365–379. MIT Press.

Denecker, M., Lierler, Y., Truszczynski, M., & Vennekens, J. (2012). A Tarskian informal semantics for ASP. In *International Conference on Logic Programming (Technical Communications)*.

Denecker, M., & Ternovska, E. (2008a). A logic of nonmonotone inductive definitions. In *ACM Transactions on Computational Logic (TOCL)* (Denecker & Ternovska, 2008b), pp. 14:1–14:52.

Denecker, M., & Ternovska, E. (2008b). A logic of nonmonotone inductive definitions. *ACM Transactions on Computational Logic (TOCL)*, *9*(2), 14:1–14:52.

Denecker, M., Vennekens, J., Bond, S., Gebser, M., & Truszczyński, M. (2009). The second Answer Set Programming competition. In *LPNMR*, pp. 637–654.

Detlefs, D., Nelson, G., & Saxe, J. B. (2005). Simplify: A theorem prover for program checking. *J. ACM*, *52*(3), 365–473.

Drescher, C., & Walsh, T. (2011). Conflict-driven constraint answer set solving with lazy nogood generation. In *AAAI*.

Eiter, T., Ianni, G., Schindlauer, R., & Tompits, H. (2005). A uniform integration of higher-order reasoning and external evaluations in answer-set programming. In Kaelbling, L. P., & Saffiotti, A. (Eds.), *IJCAI*, pp. 90–96. Professional Book Center.

Enderton, H. B. (2001). *A Mathematical Introduction To Logic* (Second edition). Academic Press.

Erdem, E., Lin, F., & Schaub, T. (Eds.). (2009). *Logic Programming and Nonmonotonic Reasoning, 10th International Conference, LPNMR 2009, Potsdam, Germany, September 14-18, 2009. Proceedings*, Vol. 5753 of *LNCS*. Springer.

Ge, Y., & de Moura, L. M. (2009). Complete instantiation for quantified formulas in satisfiabiliby modulo theories. In Bouajjani, A., & Maler, O. (Eds.), *Computer Aided Verification*, Vol. 5643 of *Lecture Notes in Computer Science*, pp. 306–320. Springer Berlin Heidelberg.

Gebser, M., Kaminski, R., Kaufmann, B., Ostrowski, M., Schaub, T., & Thiele, S. (2008). Engineering an incremental ASP solver. In García de la Banda, M., & Pontelli, E. (Eds.), *ICLP*, Vol. 5366 of *LNCS*, pp. 190–205. Springer.

Gebser, M., Kaminski, R., König, A., & Schaub, T. (2011). Advances in *gringo* series 3.. In Delgrande, & Faber (Delgrande & Faber, 2011), pp. 345–351.

Gebser, M., Schaub, T., & Thiele, S. (2007). GrinGo : A new grounder for answer set programming. In Baral, C., Brewka, G., & Schlipf, J. S. (Eds.), *LPNMR*, Vol. 4483 of *LNCS*, pp. 266–271. Springer.

Hähnle, R. (2001). Tableaux and related methods. In Robinson, J. A., & Voronkov, A. (Eds.), *Handbook of Automated Reasoning*, pp. 100–178. Elsevier and MIT Press.

Jackson, E. K., Bjorner, N., & Schulte, W. (2013). Open-world logic programs: A new foundation for formal specifications. *Tech. Report*, 1–19.

Jansen, J., Jorissen, A., & Janssens, G. (2013). Compiling input$*$ FO$(\cdot)$ inductive definitions into tabled Prolog rules for IDP^3. *Theory and Practice of Logic Programming (TPLP)*, *13*(4-5), 691–704.

Karp, R. (1972). Reducibility among combinatorial problems. In Miller, R., & Thatcher, J. (Eds.), *Complexity of Computer Computations*, pp. 85–103. Plenum Press.

Kimmig, A., Demoen, B., De Raedt, L., Santos Costa, V., & Rocha, R. (2011). On the implementation of the probabilistic logic programming language ProbLog. *TPLP*, *11*(2-3), 235–262.

Lefèvre, C., & Nicolas, P. (2009). The first version of a new ASP solver : ASPeRiX.. In Erdem et al. (Erdem, Lin, & Schaub, 2009), pp. 522–527.

Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S., & Scarcello, F. (2006). The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic (TOCL)*, *7*(3), 499–562.

Liu, G., Janhunen, T., & Niemelä, I. (2012). Answer set programming via mixed integer programming. In Brewka, G., Eiter, T., & McIlraith, S. A. (Eds.), *KR*. AAAI Press.

Marek, V. W., & Truszczyński, M. (1999). Stable models and an alternative logic programming paradigm. In Apt, K. R., Marek, V. W., Truszczyński, M., & Warren, D. S. (Eds.), *The Logic Programming Paradigm: a 25-Year Perspective*, pp. 375–398. Springer-Verlag.

Mariën, M. (2009). *Model Generation for ID-Logic*. Ph.D. thesis, Department of Computer Science, K.U.Leuven, Leuven, Belgium.

Mariën, M., Gilis, D., & Denecker, M. (2004). On the relation between ID-Logic and Answer Set Programming. In Alferes, J. J., & Leite, J. A. (Eds.), *JELIA*, Vol. 3229 of *LNCS*, pp. 108–120. Springer.

Mariën, M., Wittocx, J., Denecker, M., & Bruynooghe, M. (2008). SAT(ID): Satisfiability of propositional logic extended with inductive definitions. In *SAT*, pp. 211–224.

Metodi, A., & Codish, M. (2012). Compiling finite domain constraints to SAT with BEE. *TPLP*, *12*(4-5), 465–483.

Mitchell, D. G. (2005). A SAT solver primer. *Bulletin of the European Association for Theoretical Computer Science*, *85*, 112–132.

Mitchell, D. G., Ternovska, E., Hach, F., & Mohebali, R. (2006). Model expansion as a framework for modelling and solving search problems. Tech. rep. TR 2006-24, Simon Fraser University, Canada.

Moskewicz, M., Madigan, C., Zhao, Y., Zhang, L., & Malik, S. (2001). Chaff: Engineering an efficient SAT solver. In *DAC'01*, pp. 530–535. ACM.

Nethercote, N., Stuckey, P., Becket, R., Brand, S., Duck, G., & Tack, G. (2007). Minizinc: Towards a standard CP modelling language. In Bessiere, C. (Ed.), *CP'07*, Vol. 4741 of *LNCS*, pp. 529–543. Springer.

Nightingale, P., Gent, I. P., Jefferson, C., & Miguel, I. (2013). Short and long supports for constraint propagation. *J. Artif. Intell. Res. (JAIR)*, *46*, 1–45.

Ostrowski, M., & Schaub, T. (2012). ASP modulo CSP: The clingcon system. *TPLP*, *12*(4-5), 485–503.

Palù, A. D., Dovier, A., Pontelli, E., & Rossi, G. (2009). Answer set programming with constraints using lazy grounding. In Hill, P. M., & Warren, D. S. (Eds.), *ICLP*, Vol. 5649 of *LNCS*, pp. 115–129. Springer.

Riedel, S. (2009). Cutting plane map inference for markov logic. In *SRL 2009*.

Saptawijaya, A., & Pereira, L. M. (2013). Towards practical tabled abduction in logic programs. In Correia, L., Reis, L. P., & Cascalho, J. (Eds.), *EPIA*, Vol. 8154 of *LNCS*, pp. 223–234. Springer.

Singla, P., & Domingos, P. (2006). Memory-efficient inference in relational domains. In Gil, Y., & Mooney, R. J. (Eds.), *AAAI*, pp. 488–493. AAAI Press.

Stuckey, P. J. (2010). Lazy clause generation: Combining the power of SAT and CP (and MIP?) solving. In *CPAIOR*, pp. 5–9.

Tamura, N., Taga, A., Kitagawa, S., & Banbara, M. (2009). Compiling finite linear CSP into SAT. *Constraints*, *14*(2), 254–272.

Torlak, E., Chang, F. S.-H., & Jackson, D. (2008). Finding minimal unsatisfiable cores of declarative specifications. In Cuéllar, J., Maibaum, T. S. E., & Sere, K. (Eds.), *FM*, Vol. 5014 of *LNCS*, pp. 326–341. Springer.

Van Gelder, A. (1993). The alternating fixpoint of logic programs with negation. *Journal of Computer and System Sciences*, *47*(1), 185–221.

Vennekens, J., Mariën, M., Wittocx, J., & Denecker, M. (2007). Predicate introduction for logics with a fixpoint semantics. Part I: Logic programming. *Fundamenta Informaticae*, *79*(1-2), 187–208.

Wittocx, J., Denecker, M., & Bruynooghe, M. (2013a). Constraint propagation for first-order logic and inductive definitions. In *ACM Trans. Comput. Logic* (Wittocx et al., 2013b), pp. 17:1–17:45.

Wittocx, J., Denecker, M., & Bruynooghe, M. (2013b). Constraint propagation for first-order logic and inductive definitions. *ACM Trans. Comput. Logic*, *14*(3), 17:1–17:45.

Wittocx, J., Mariën, M., & Denecker, M. (2008). The IDP system: a model expansion system for an extension of classical logic. In Denecker, M. (Ed.), *LaSh*, pp. 153–165.

Wittocx, J., Mariën, M., & Denecker, M. (2010a). Grounding FO and FO(ID) with bounds. In *Journal of Artificial Intelligence Research* (Wittocx et al., 2010b), pp. 223–269.

Wittocx, J., Mariën, M., & Denecker, M. (2010b). Grounding FO and FO(ID) with bounds. *Journal of Artificial Intelligence Research*, *38*, 223–269.